SAND2014-5015
Unlimited Release
July 2014
Updated August 5, 2014

Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.0 Reference Manual

Brian M. Adams, Mohamed S. Ebeida, Michael S. Eldred, John D. Jakeman, Laura P. Swiler, J. Adam Stephens, Dena M. Vigil, Timothy M. Wildey
Optimization and Uncertainty Quantification Department

William J. Bohnhoff
Radiation Transport Department

Keith R. Dalbey
Mission Analysis and Simulation Department

John P. Eddy
System Readiness and Sustainment Technologies Department

Kenneth T. Hu
Validation and Uncertainty Quantification Department

Lara E. Bauman, Patricia D. Hough
Quantitative Modeling and Analysis Department

Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185

Abstract

The Dakota (Design Analysis Kit for Optimization and Terascale Applications) toolkit provides a flexible and extensible interface between simulation codes and iterative analysis methods. Dakota contains algorithms for optimization with gradient and nongradient-based methods; uncertainty quantification with sampling, reliability, and stochastic expansion methods; parameter estimation with nonlinear least squares methods; and sensitivity/variance analysis with design of experiments and parameter study methods. These capabilities may be used on their own or as components within advanced strategies such as surrogate-based optimization, mixed integer nonlinear programming, or optimization under uncertainty. By employing object-oriented design to implement abstractions of the key components required for iterative systems analyses, the Dakota toolkit provides a flexible and extensible problem-solving environment for design and performance analysis of computational models on high performance computers.

This report serves as a reference manual for the commands specification for the Dakota software, providing input overviews, option descriptions, and example specifications.

Contents

1 Main Page **7**
1.1 How to Use this Manual . . . 7

2 Running Dakota **9**
2.1 Usage . . . 9
2.2 Examples . . . 10
2.3 Execution Phases . . . 10
2.4 Restarting Dakota Studies . . . 11
2.5 The Dakota Restart Utility . . . 12

3 Dakota Input Specification **17**
3.1 Dakota NIDR . . . 17
3.2 Input Spec Overview . . . 17
3.3 Sample Input Files . . . 19
3.4 Input Spec Summary . . . 23

4 Topics Area **59**
4.1 admin . . . 59
4.2 dakota_IO . . . 70
4.3 dakota_concepts . . . 77
4.4 models . . . 95
4.5 variables_and_responses . . . 97
4.6 interface . . . 103
4.7 methods . . . 105
4.8 advanced_topics . . . 125
4.9 packages . . . 129

5 Keywords Area **137**
5.1 environment . . . 138
5.2 method . . . 143
5.3 model . . . 1137
5.4 variables . . . 1226
5.5 interface . . . 1337
5.6 responses . . . 1390

Bibliographic References **1453**

Chapter 1

Main Page

The **Dakota** (**D**esign **A**nalysis **K**it for **O**ptimization and **T**erascale **A**pplications) toolkit provides a flexible, extensible interface between analysis codes and iteration methods.

Author

> Brian M. Adams, Lara E. Bauman, William J. Bohnhoff, Keith R. Dalbey, John P. Eddy, Mohamed S. Ebeida, Michael S. Eldred, Patricia D. Hough, Kenneth T. Hu, John D. Jakeman, Laura P. Swiler, J. Adam Stephens, Dena M. Vigil, Timothy M. Wildey

The Reference Manual documents all the input keywords that can appear in a Dakota input file to configure a Dakota study. Its organization closely mirrors the structure of `dakota.input.summary`. For more information see Dakota Input Specification. For information on software structure, refer to the Developers Manual [3], and for a tour of Dakota features and capabilities, including a tutorial, refer to the User's Manual (Adams et al., 2010) [4].

1.1 How to Use this Manual

- To learn how to run Dakota from the command line, see Running Dakota
- To learn to how to restart Dakota studies, see Restarting Dakota Studies
- To learn about the Dakota restart utility, see The Dakota Restart Utility

To find more information about a specific keyword

1. Use the search box at the top right (currently only finds keyword names)
2. Browse the Keywords tree on the left navigation pane
3. Look at the Dakota Input Specification
4. Navigate through the keyword pages, starting from the Keywords Area

To find more information about a Dakota related topic

1. Browse the Topics Area on the left navigation pane
2. Navigate through the topics pages, starting from the Topics Area

A bibliography for the Reference Manual is provided in Bibliographic References

Chapter 2

Running Dakota

The Dakota executable file is named `dakota` (`dakota.exe` on Windows) and is most commonly run from a terminal or command prompt.

2.1 Usage

If the `dakota` command is entered at the command prompt without any arguments, a usage message similar to the following appears:

```
usage: dakota [options and <args>]
        -help (Print this summary)
        -version (Print Dakota version number)
        -input <$val> (REQUIRED Dakota input file $val)
        -output <$val> (Redirect Dakota standard output to file $val)
        -error <$val> (Redirect Dakota standard error to file $val)
        -parser <$val> (Parsing technology: nidr[strict][:dumpfile])
        -no_input_echo (Do not echo Dakota input file)
        -check (Perform input checks)
        -pre_run [$val] (Perform pre-run (variables generation) phase)
        -run [$val] (Perform run (model evaluation) phase)
        -post_run [$val] (Perform post-run (final results) phase)
        -read_restart [$val] (Read an existing Dakota restart file $val)
        -stop_restart <$val> (Stop restart file processing at evaluation $val)
        -write_restart [$val] (Write a new Dakota restart file $val)
```

Of these command line options, only `input` is required, and the `-input` switch can be omitted if the input file name is the final item appearing on the command line (see Examples); all other command-line inputs are optional.

- `help` prints the usage message above.
- `version` prints version information for the executable.
- `check` invokes a dry-run mode in which the input file is processed and checked for errors, but the study is not performed.
- `input` provides the name of the Dakota input file.
- `output` and `error` options provide file names for redirection of the Dakota standard output (stdout) and standard error (stderr), respectively.

- The `parser` option is for debugging and will not be further described here.
- By default, Dakota will echo the input file to the output stream, but `no_input_echo` can override this behavior.
- `read_restart` and `write_restart` commands provide the names of restart databases to read from and write to, respectively.
- `stop_restart` command limits the number of function evaluations read from the restart database (the default is all the evaluations) for those cases in which some evaluations were erroneous or corrupted. Restart management is an important technique for retaining data from expensive engineering applications.
- `-pre_run`, `-run`, and `-post_run` instruct Dakota to run one or more execution phases, excluding others. The commands must be followed by filenames as described in Execution Phases.

Command line switches can be abbreviated so long as the abbreviation is unique, so the following are valid, unambiguous specifications: `-h`, `-v`, `-c`, `-i`, `-o`, `-e`, `-s`, `-w`, `-re`, `-pr`, `-ru`, and `-po` and can be used in place of the longer forms of the command line options.

For information on restarting Dakota, see Restarting Dakota Studies and The Dakota Restart Utility.

2.2 Examples

To run Dakota with a particular input file, the following syntax can be used:

```
dakota -i dakota.in
```

or more simply

```
dakota dakota.in
```

This will echo the standard output (stdout) and standard error (stderr) messages to the terminal. To redirect stdout and stderr to separate files, the `-o` and `-e` command line options may be used:

```
dakota -i dakota.in -o dakota.out -e dakota.err
```

or

```
dakota -o dakota.out -e dakota.err dakota.in
```

Alternatively, any of a variety of Unix redirection variants can be used. Refer to[6] for more information on Unix redirection. The simplest of these redirects stdout to another file:

```
dakota dakota.in > dakota.out
```

2.3 Execution Phases

Dakota has three execution phases: `pre-run`, `run`, and `post-run`.

- `pre-run` can be used to generate variable sets
- `run` (core run) invokes the simulation to evaluate variables, producing responses
- `post-run` accepts variable/response sets and analyzes the results (for example, calculate correlations from a set of samples). Currently only two modes are supported and only for sampling, parameter study, and DACE methods:

(1) pre-run only with optional tabular output of variables:

```
dakota -i dakota.in -pre_run [::myvariables.dat]
```

(2) post-run only with required tabular input of variables/responses:

```
dakota -i dakota.in -post_run myvarsresponses.dat::
```

2.4 Restarting Dakota Studies

Dakota is often used to solve problems that require repeatedly running computationally expensive simulation codes. In some cases you may want to repeat an optimization study, but with a tighter final convergence tolerance. This would be costly if the entire optimization analysis had to be repeated. Interruptions imposed by computer usage policies, power outages, and system failures could also result in costly delays. However, Dakota automatically records the variable and response data from all function evaluations so that new executions of Dakota can pick up where previous executions left off. The Dakota restart file (`dakota.rst` by default) archives the tabulated interface evaluations in a binary format. The primary restart commands at the command line are `-read_restart`, `-write_restart`, and `-stop_restart`.

2.4.1 Writing Restart Files

To write a restart file using a particular name, the `-write_restart` command line input (may be abbreviated as `-w`) is used:

```
dakota -i dakota.in -write_restart my_restart_file
```

If no `-write_restart` specification is used, then Dakota will still write a restart file, but using the default name `dakota.rst` instead of a user-specified name.

To turn restart recording off, the user may use the restart_file keyword, in the `interface` block. This can increase execution speed and reduce disk storage requirements, but at the expense of a loss in the ability to recover and continue a run that terminates prematurely. This option is not recommended when function evaluations are costly or prone to failure. Please note that using the `deactivate restart_file` specification will result in a zero length restart file with the default name dakota.rst, which can overwrite an exiting file.

2.4.2 Using Restart Files

To restart Dakota from a restart file, the `-read_restart` command line input (may be abbreviated as `-r`) is used:

```
dakota -i dakota.in -read_restart my_restart_file
```

If no `-read_restart` specification is used, then Dakota will not read restart information from any file (i.e., the default is no restart processing).

To read in only a portion of a restart file, the `-stop_restart` control (may be abbreviated as `-s`) is used to specify the number of entries to be read from the database. Note that this integer value corresponds to the restart record processing counter (as can be seen when using the `print` utility (see The Dakota Restart Utility) which may differ from the evaluation numbers used in the previous run if, for example, any duplicates were detected (since these duplicates are not recorded in the restart file). In the case of a `-stop_restart` specification, it is usually desirable to specify a new restart file using `-write_restart` so as to remove the records of erroneous or corrupted function evaluations. For example, to read in the first 50 evaluations from `dakota.rst`:

```
dakota -i dakota.in -r dakota.rst -s 50 -w dakota_new.rst
```

The `dakota_new.rst` file will contain the 50 processed evaluations from `dakota.rst` as well as any new evaluations. All evaluations following the 50th in `dakota.rst` have been removed from the latest restart record.

2.4.3 Appending to a Restart File

If the `-write_restart` and `-read_restart` specifications identify the same file (including the case where `-write_restart` is not specified and `-read_restart` identifies `dakota.rst`), then new evaluations will be appended to the existing restart file.

2.4.4 Working with multiple Restart Files

If the `-write_restart` and `-read_restart` specifications identify different files, then the evaluations read from the file identified by `-read_restart` are first written to the `-write_restart` file. Any new evaluations are then appended to the `-write_restart` file. In this way, restart operations can be chained together indefinitely with the assurance that all of the relevant evaluations are present in the latest restart file.

2.4.5 How it Works

Dakota's restart algorithm relies on its duplicate detection capabilities. Processing a restart file populates the list of function evaluations that have been performed. Then, when the study is restarted, it is started from the beginning (not a warm start) and many of the function evaluations requested by the iterator are intercepted by the duplicate detection code. This approach has the primary advantage of restoring the complete state of the iteration (including the ability to correctly detect subsequent duplicates) for all methods/iterators without the need for iterator-specific restart code. However, the possibility exists for numerical round-off error to cause a divergence between the evaluations performed in the previous and restarted studies. This has been rare in practice.

2.5 The Dakota Restart Utility

The Dakota restart utility program provides a variety of facilities for managing restart files from Dakota executions. The executable program name is `dakota_restart_util` and it has the following options, as shown by the usage message returned when executing the utility without any options:

```
Usage:
  dakota_restart_util print <restart_file>
  dakota_restart_util to_neutral <restart_file> <neutral_file>
  dakota_restart_util from_neutral <neutral_file> <restart_file>
  dakota_restart_util to_pdb <restart_file> <pdb_file>
  dakota_restart_util to_tabular <restart_file> <text_file>
  dakota_restart_util remove <double> <old_restart_file> <new_restart_file>
  dakota_restart_util remove_ids <int_1> ... <int_n> <old_restart_file> <new_restart_file>
  dakota_restart_util cat <restart_file_1> ... <restart_file_n> <new_restart_file>
```

Several of these functions involve format conversions. In particular, the binary format used for restart files can be converted to ASCII text and printed to the screen, converted to and from a neutral file format, converted to a PDB format for use at Lawrence Livermore National Laboratory, or converted to a tabular format for importing into 3rd-party graphics programs. In addition, a restart file with corrupted data can be repaired by value or id, and multiple restart files can be combined to create a master database.

2.5.1 Print Command

The `print` option is useful to show contents of a restart file, since the binary format is not convenient for direct inspection. The restart data is printed in full precision, so that exact matching of points is possible for restarted runs or corrupted data removals. For example, the following command

```
dakota_restart_util print
dakota.rst
```

results in output similar to the following:

```
----------------------------------------
Restart record    1  (evaluation id    1):
----------------------------------------
Parameters:
                      1.8000000000000000e+00 intake_dia
                      1.0000000000000000e+00 flatness
```

```
Active response data:
Active set vector = { 3 3 3 3 }
                    -2.4355973813420619e+00 obj_fn
                    -4.7428486677140930e-01 nln_ineq_con_1
                    -4.5000000000000001e-01 nln_ineq_con_2
                     1.3971143170299741e-01 nln_ineq_con_3
 [ -4.3644298963447897e-01  1.4999999999999999e-01 ] obj_fn gradient
 [  1.3855136437818300e-01  0.0000000000000000e+00 ] nln_ineq_con_1 gradient
 [  0.0000000000000000e+00  1.4999999999999999e-01 ] nln_ineq_con_2 gradient
 [  0.0000000000000000e+00 -1.9485571585149869e-01 ] nln_ineq_con_3 gradient

----------------------------------------
Restart record    2  (evaluation id    2):
----------------------------------------
Parameters:
                     2.1640000000000001e+00 intake_dia
                     1.7169994018008317e+00 flatness

Active response data:
Active set vector = { 3 3 3 3 }
                    -2.4869127192988878e+00 obj_fn
                     6.9256958799989843e-01 nln_ineq_con_1
                    -3.4245008972987528e-01 nln_ineq_con_2
                     8.7142207937157910e-03 nln_ineq_con_3
 [ -4.3644298963447897e-01  1.4999999999999999e-01 ] obj_fn gradient
 [  2.9814239699997572e+01  0.0000000000000000e+00 ] nln_ineq_con_1 gradient
 [  0.0000000000000000e+00  1.4999999999999999e-01 ] nln_ineq_con_2 gradient
 [  0.0000000000000000e+00 -1.6998301774282701e-01 ] nln_ineq_con_3 gradient

...<snip>...

Restart file processing completed: 11 evaluations retrieved.
```

2.5.2 Neutral File Format

A Dakota restart file can be converted to a neutral file format using a command like the following:

```
dakota_restart_util to_neutral dakota.rst dakota.neu
```

which results in a report similar to the following:

```
Writing neutral file dakota.neu
Restart file processing completed: 11 evaluations retrieved.
```

Similarly, a neutral file can be returned to binary format using a command like the following:

```
dakota_restart_util from_neutral dakota.neu dakota.rst
```

which results in a report similar to the following:

```
Reading neutral file dakota.neu
Writing new restart file dakota.rst
Neutral file processing completed: 11 evaluations retrieved.
```

The contents of the generated neutral file are similar to the following (from the first two records for the Cylinder example in[4]).

```
6 7 2 1.8000000000000000e+00 intake_dia 1.0000000000000000e+00 flatness 0 0 0 0
NULL 4 2 1 0 3 3 3 3 1 2 obj_fn nln_ineq_con_1 nln_ineq_con_2 nln_ineq_con_3
  -2.4355973813420619e+00 -4.7428486677140930e-01 -4.5000000000000001e-01
   1.3971143170299741e-01 -4.3644298963447897e-01  1.4999999999999999e-01
   1.3855136437818300e-01  0.0000000000000000e+00  0.0000000000000000e+00
   1.4999999999999999e-01  0.0000000000000000e+00 -1.9485571585149869e-01 1
```

```
6 7 2 2.1640000000000001e+00 intake_dia 1.7169994018008317e+00 flatness 0 0 0 0
NULL 4 2 1 0 3 3 3 3 1 2 obj_fn nln_ineq_con_1 nln_ineq_con_2 nln_ineq_con_3
  -2.4869127192988878e+00 6.9256958799989843e-01 -3.4245008972987528e-01
   8.7142207937157910e-03 -4.3644298963447897e-01  1.4999999999999999e-01
   2.9814239699997572e+01  0.0000000000000000e+00  0.0000000000000000e+00
   1.4999999999999999e-01  0.0000000000000000e+00 -1.6998301774282701e-01 2
```

This format is not intended for direct viewing (`print` should be used for this purpose). Rather, the neutral file capability has been used in the past for managing portability of restart data across platforms (recent use of more portable binary formats has largely eliminated this need) or for advanced repair of restart records (in cases where the `remove` command was insufficient).

2.5.3 Tabular Format

Conversion of a binary restart file to a tabular format enables convenient import of this data into 3rd-party post-processing tools such as Matlab, TECplot, Excel, etc. This facility is nearly identical to the output activated by the tabular_graphics_data keyword in the Dakota input file specification, but with two important differences:

1. No function evaluations are suppressed as they are with `tabular_graphics_data` (i.e., any internal finite difference evaluations are included).
2. The conversion can be performed later, i.e., for Dakota runs executed previously.

An example command for converting a restart file to tabular format is:

```
dakota_restart_util to_tabular dakota.rst dakota.m
```

which results in a report similar to the following:

```
Writing tabular text file dakota.m
Restart file processing completed: 10 evaluations tabulated.
```

The contents of the generated tabular file are similar to the following (from the example in[4]). Note that while evaluations resulting from numerical derivative offsets would be reported (as described above), derivatives returned as part of the evaluations are not reported (since they do not readily fit within a compact tabular format):

```
%eval_id             x1             x2         obj_fn nln_ineq_con_1 nln_ineq_con_2
       1            0.9            1.1         0.0002           0.26           0.76
       2   0.6433962264   0.6962264151   0.0246865569  0.06584549662   0.1630331079
       3   0.5310576935   0.5388046558  0.09360081618  0.01261994597  0.02478161031
       4    0.612538853   0.6529854907  0.03703861037  0.04871110113   0.1201206246
       5   0.5209215947   0.5259311717   0.1031862798  0.00839372202  0.01614279999
       6   0.5661606434   0.5886684401  0.06405197568  0.02620365411  0.06345021064
       7   0.5083873357    0.510239856   0.1159458957 0.003337755086 0.006151042802
       8   0.5001577143   0.5001800249   0.1248312163 6.772666885e-05 0.0001012002012
       9   0.5000000547   0.5000000598   0.1249999428 2.485652461e-08 3.238746073e-08
      10            0.5            0.5          0.125 2.942091015e-15 3.60822483e-15
```

2.5.4 Concatenation of Multiple Restart Files

In some instances, it is useful to combine restart files into a single master function evaluation database. For example, when constructing a data fit surrogate model, data from previous studies can be pulled in and reused to create a combined data set for the surrogate fit. An example command for concatenating multiple restart files is:

```
dakota_restart_util cat dakota.rst.1 dakota.rst.2 dakota.rst.3 dakota.rst.all
```

which results in a report similar to the following:

```
Writing new restart file dakota.rst.all
dakota.rst.1 processing completed: 10 evaluations retrieved.
dakota.rst.2 processing completed: 110 evaluations retrieved.
dakota.rst.3 processing completed: 65 evaluations retrieved.
```

The `dakota.rst.all` database now contains 185 evaluations and can be read in for use in a subsequent Dakota study using the `-read_restart` option to the `dakota` executable.

2.5.5 Removal of Corrupted Data

On occasion, a simulation or computer system failure may cause a corruption of the Dakota restart file. For example, a simulation crash may result in failure of a post-processor to retrieve meaningful data. If 0's (or other erroneous data) are returned from the user's `analysis_driver`, then this bad data will get recorded in the restart file. If there is a clear demarcation of where corruption initiated (typical in a process with feedback, such as gradient-based optimization), then use of the `-stop_restart` option for the `dakota` executable can be effective in continuing the study from the point immediately prior to the introduction of bad data. If, however, there are interspersed corruptions throughout the restart database (typical in a process without feedback, such as sampling), then the `remove` and `remove_ids` options of `dakota_restart_util` can be useful.

An example of the command syntax for the `remove` option is:

```
dakota_restart_util remove 2.e-04 dakota.rst dakota.rst.repaired
```

which results in a report similar to the following:

```
Writing new restart file dakota.rst.repaired
Restart repair completed: 65 evaluations retrieved, 2 removed, 63 saved.
```

where any evaluations in `dakota.rst` having an active response function value that matches `2.e-04` within machine precision are discarded when creating `dakota.rst.repaired`.

An example of the command syntax for the `remove_ids` option is:

```
dakota_restart_util remove_ids 12 15 23 44 57 dakota.rst dakota.rst.repaired
```

which results in a report similar to the following:

```
Writing new restart file dakota.rst.repaired
Restart repair completed: 65 evaluations retrieved, 5 removed, 60 saved.
```

where evaluation ids `12`, `15`, `23`, `44`, and `57` have been discarded when creating `dakota.rst.repaired`. An important detail is that, unlike the `-stop_restart` option which operates on restart record numbers, the `remove_ids` option operates on evaluation ids. Thus, removal is not necessarily based on the order of appearance in the restart file. This distinction is important when removing restart records for a run that contained either asynchronous or duplicate evaluations, since the restart insertion order and evaluation ids may not correspond in these cases (asynchronous evaluations have ids assigned in the order of job creation but are inserted in the restart file in the order of job completion, and duplicate evaluations are not recorded which introduces offsets between evaluation id and record number). This can also be important if removing records from a concatenated restart file, since the same evaluation id could appear more than once. In this case, all evaluation records with ids matching the `remove_ids` list will be removed.

If neither of these removal options is sufficient to handle a particular restart repair need, then the fallback position is to resort to direct editing of a neutral file to perform the necessary modifications.

Chapter 3

Dakota Input Specification

3.1 Dakota NIDR

Valid Dakota input is dictated governed by the NIDR[29] input specification file, `dakota.input.nspec`. This file is used by a code generator to create parsing system components that are compiled into Dakota. Therefore, dakota.input.nspec and its derived summary, dakota.input.summary, are the definitive source for input syntax, capability options, and optional and required capability sub-parameters for any given Dakota version.

Beginning users may find dakota.input.summary overwhelming or confusing and will likely derive more benefit from adapting example input files to a particular problem. Some examples can be found here: Sample Input Files. Advanced users can master the many input specification possibilities by understanding the structure of the input specification file.

3.2 Input Spec Overview

Refer to the dakota.input.summary file, in Input Spec Summary, for current input specifications.

- The summary describes every keyword including:
  - Whether it is required or optional
  - Whether it takes `ARGUMENTS` (always required) Additional notes about `ARGUMENTS` can be found here: Specifying Arguments.
  - Whether it has an `ALIAS`, or synonym
  - Which additional keywords can be specified to change its behavior
- Additional details and descriptions are described in Keywords Area
- For additional details on NIDR specification logic and rules, refer to[29] (Gay, 2008).

3.2.1 Common Specification Mistakes

Spelling mistakes and omission of required parameters are the most common errors. Some causes of errors are more obscure:

- Documentation of new capability sometimes lags its availability in source and executables, especially stable releases. When parsing errors occur that the documentation cannot explain, reference to the particular input specification used in building the executable, which is installed alongside the executable, will often resolve the errors.

- If you want to compare results with those obtained using an earlier version of Dakota (prior to 4.1), your input file for the earlier version must use backslashes to indicate continuation lines for Dakota keywords. For example, rather than

  ```
  # Comment about the following "responses" keyword...
  responses,
      objective_functions = 1
      # Comment within keyword "responses"
      analytic_gradients
  # Another comment within keyword "responses"
      no_hessians
  ```

  you would need to write

  ```
  # Comment about the following "responses" keyword...
  responses,                        \
      objective_functions = 1           \
      # Comment within keyword "responses"      \
      analytic_gradients              \
  # Another comment within keyword "responses"      \
      no_hessians
  ```

  with no white space (blanks or tabs) after the \ character.

In most cases, the NIDR system provides error messages that help the user isolate errors in Dakota input files.

3.2.2 Specifying Arguments

Some keywords, such as those providing bounds on variables, have an associated list of values or strings, referred to as arguments.

When the same value should be repeated several times in a row, you can use the notation `N*value` instead of repeating the value N times.

For example

```
lower_bounds  -2.0  -2.0  -2.0
upper_bounds   2.0   2.0   2.0
```

could also be written

```
lower_bounds  3*-2.0
upper_bounds  3* 2.0
```

(with optional spaces around the $*$).

Another possible abbreviation is for sequences: L:S:U (with optional spaces around the :) is expanded to L L+S L+2$*$S ... U, and L:U (with no second colon) is treated as L:1:U.

For example, in one of the test examples distributed with Dakota (test case 2 of `test/dakota_uq_textbook_sop_lhs.in`),

```
histogram_point = 2
  abscissas   = 50. 60. 70. 80. 90.
                30. 40. 50. 60. 70.
  counts      = 10 20 30 20 10
                10 20 30 20 10
```

could also be written

```
histogram_point = 2
  abscissas   = 50 : 10 : 90
                30 : 10 : 70
  counts      = 10:10:30 20 10
                10:10:30 20 10
```

Count and sequence abbreviations can be used together. For example

```
response_levels =
    0.0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1.0
    0.0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1.0
```

can be abbreviated

```
response_levels =
    2*0.0:0.1:1.0
```

3.3 Sample Input Files

A Dakota input file is a collection of fields from the dakota.input.summary file that describe the problem to be solved by Dakota. Several examples follow.

Sample 1: Optimization

The following sample input file shows single-method optimization of the Textbook Example using DOT's modified method of feasible directions. A similar file is available as `Dakota/examples/users/textbook_opt_conmin.in`.

```
# Dakota Input File: textbook_opt_conmin.in
environment
 graphics
 tabular_graphics_data
  tabular_graphics_file = 'textbook_opt_conmin.dat'

method
# dot_mmfd #DOT performs better but may not be available
 conmin_mfd
  max_iterations = 50
  convergence_tolerance = 1e-4

variables
 continuous_design = 2
  initial_point  0.9  1.1
  upper_bounds   5.8  2.9
  lower_bounds   0.5  -2.9
  descriptors   'x1'  'x2'

interface
 direct
  analysis_driver =    'text_book'

responses
 objective_functions = 1
 nonlinear_inequality_constraints = 2
 numerical_gradients
  method_source dakota
  interval_type central
  fd_gradient_step_size = 1.e-4
 no_hessians
```

Sample 2: Least Squares (Calibration)

The following sample input file shows a nonlinear least squares (calibration) solution of the Rosenbrock Example using the NL2SOL method. A similar file is available as `Dakota/examples/users/rosen_opt_nls.in`

```
# Dakota Input File: rosen_opt_nls.in
environment
```

```
 graphics
 tabular_graphics_data
  tabular_graphics_file = 'rosen_opt_nls.dat'

method
 max_iterations = 100
 convergence_tolerance = 1e-4
 nl2sol

model
 single

variables
 continuous_design = 2
  initial_point  -1.2   1.0
  lower_bounds   -2.0   -2.0
  upper_bounds   2.0   2.0
  descriptors    'x1'   "x2"

interface
 analysis_driver = 'rosenbrock'
  direct

responses
 calibration_terms = 2
 analytic_gradients
 no_hessians
```

Sample 3: Nondeterministic Analysis

The following sample input file shows Latin Hypercube Monte Carlo sampling using the Textbook Example. A similar file is available as `Dakota/test/dakota_uq_textbook_lhs.in`.

```
method,
    sampling,
     samples = 100 seed = 1
     complementary distribution
     response_levels = 3.6e+11 4.0e+11 4.4e+11
              6.0e+04 6.5e+04 7.0e+04
              3.5e+05 4.0e+05 4.5e+05
     sample_type lhs

variables,
    normal_uncertain = 2
     means       = 248.89, 593.33
     std_deviations  =  12.4,  29.7
     descriptors    = 'TF1n' 'TF2n'
    uniform_uncertain = 2
     lower_bounds   = 199.3, 474.63
     upper_bounds   = 298.5, 712.
     descriptors    = 'TF1u' 'TF2u'
    weibull_uncertain = 2
     alphas      =  12.,  30.
     betas       = 250.,  590.
     descriptors    = 'TF1w' 'TF2w'
    histogram_bin_uncertain = 2
     num_pairs  = 3     4
     abscissas  = 5 8 10 .1 .2 .3 .4
     counts   = 17 21 0 12 24 12  0
     descriptors = 'TF1h' 'TF2h'
    histogram_point_uncertain = 1
     num_pairs  = 2
     abscissas  = 3 4
```

```
     counts   = 1 1
     descriptors = 'TF3h'

interface,
    fork asynch evaluation_concurrency = 5
     analysis_driver = 'text_book'

responses,
    response_functions = 3
    no_gradients
    no_hessians
```

Sample 4: Parameter Study

The following sample input file shows a 1-D vector parameter study using the Textbook Example. It makes use of the default environment and model specifications, so they can be omitted. A similar file is available in the test directory as `Dakota/examples/users/rosen_ps_vector.in`.

```
# Dakota Input File: rosen_ps_vector.in
environment
 graphics
 tabular_graphics_data
  tabular_graphics_file = 'rosen_ps_vector.dat'

method
 vector_parameter_study
  final_point = 1.1 1.3
  num_steps = 10

variables
 continuous_design = 2
  initial_point  -0.3   0.2
  descriptors    'x1'   "x2"

interface
 analysis_driver = 'rosenbrock'
  direct

responses
 objective_functions = 1
 no_gradients
 no_hessians
```

Sample 5: Hybrid Strategy

The following sample input file shows a hybrid environment using three methods. It employs a genetic algorithm, pattern search, and full Newton gradient-based optimization in succession to solve the Textbook Example. A similar file is available as `Dakota/examples/users/textbook_hybrid_strat.in`.

```
environment
 graphics
 hybrid sequential
  method_list = 'PS' 'PS2' 'NLP'

method
 id_method = 'PS'
 model_pointer = 'M1'
 coliny_pattern_search stochastic
  seed = 1234
  initial_delta = 0.1
  threshold_delta = 1.e-4
  solution_accuracy = 1.e-10
```

```
  exploratory_moves basic_pattern
  #verbose output

method
 id_method = 'PS2'
 model_pointer = 'M1'
 max_function_evaluations = 10
 coliny_pattern_search stochastic
  seed = 1234
  initial_delta = 0.1
  threshold_delta = 1.e-4
  solution_accuracy = 1.e-10
  exploratory_moves basic_pattern
  #verbose output

method
 id_method = 'NLP'
 model_pointer = 'M2'
    optpp_newton
  gradient_tolerance = 1.e-12
  convergence_tolerance = 1.e-15
  #verbose output

model
 id_model = 'M1'
 single
  variables_pointer = 'V1'
  interface_pointer = 'I1'
  responses_pointer = 'R1'

model
 id_model = 'M2'
 single
  variables_pointer = 'V1'
  interface_pointer = 'I1'
  responses_pointer = 'R2'

variables
 id_variables = 'V1'
 continuous_design = 2
  initial_point  0.6  0.7
  upper_bounds   5.8  2.9
  lower_bounds   0.5  -2.9
  descriptors   'x1'  'x2'

interface
 id_interface = 'I1'
 direct
  analysis_driver= 'text_book'

responses
 id_responses = 'R1'
 objective_functions = 1
 no_gradients
 no_hessians

responses
 id_responses = 'R2'
 objective_functions = 1
 analytic_gradients
 analytic_hessians
```

Additional example input files, as well as the corresponding output and graphics, are provided in the Tutorial chapter of the Users Manual [4] (Adams et al., 2010).

3.4 Input Spec Summary

This file is derived automatically from dakota.input.nspec, which is used in the generation of parser system files that are compiled into the Dakota executable. Therefore, these files are the definitive source for input syntax, capability options, and associated data inputs. Refer to the Developers Manual information on how to modify the input specification and propagate the changes through the parsing system.

Key features of the input specification and the associated user input files include:

- In the input specification, required individual specifications simply appear, optional individual and group specifications are enclosed in [], required group specifications are enclosed in (), and either-or relationships are denoted by the | symbol. These symbols only appear in dakota.input.nspec; they must not appear in actual user input files.

- Keyword specifications (i.e., `environment`, `method`, `model`, `variables`, `interface`, and `responses`) begin with the keyword possibly preceded by white space (blanks, tabs, and newlines) both in the input specifications and in user input files. For readability, keyword specifications may be spread across several lines. Earlier versions of Dakota (prior to 4.1) required a backslash character (\) at the ends of intermediate lines of a keyword. While such backslashes are still accepted, they are no longer required.

- Some of the keyword components within the input specification indicate that the user must supply `INTEGER`, `REAL`, `STRING`, `INTEGERLIST`, `REALLIST`, or `STRINGLIST` data as part of the specification. In a user input file, the `"="` is optional, data in a `LIST` can be separated by commas or whitespace, and the `STRING` data are enclosed in single or double quotes (e.g., `'text_book'` or "text_book").

- In user input files, input is largely order-independent (except for entries in lists of data), case insensitive, and white-space insensitive. Although the order of input shown in the Sample Input Files generally follows the order of options in the input specification, this is not required.

- In user input files, specifications may be abbreviated so long as the abbreviation is unique. For example, the `npsol_sqp` specification within the method keyword could be abbreviated as `npsol`, but `dot_sqp` should not be abbreviated as `dot` since this would be ambiguous with other DOT method specifications.

- In both the input specification and user input files, comments are preceded by `#`.

- ALIAS refers to synonymous keywords, which often exist for backwards compatability. Users are encouraged to use the most current keyword.

```
KEYWORD01 environment
    [ graphics ]
    [ tabular_graphics_data
      [ tabular_graphics_file STRING ]
      ]
    [ output_precision INTEGER >= 0 ]
    [ results_output
      [ results_output_file STRING ]
      ]
    [ top_method_pointer ALIAS method_pointer STRING ]

KEYWORD12 method
    [ id_method STRING ]
    [ output
```

```
    debug
    | verbose
    | normal
    | quiet
    | silent
    ]
  [ max_iterations INTEGER >= 0 ]
  [ max_function_evaluations INTEGER >= 0 ]
  [ speculative ]
  [ convergence_tolerance REAL ]
  [ constraint_tolerance REAL ]
  [ scaling ]
  [ final_solutions INTEGER >= 0 ]
  ( hybrid
    ( sequential ALIAS uncoupled
      ( method_name_list STRINGLIST
        [ model_pointer_list STRING ]
        )
      | method_pointer_list STRINGLIST
      )
    |
    ( embedded ALIAS coupled
      ( global_method_name STRING
        [ global_model_pointer STRING ]
        )
      | global_method_pointer STRING
      ( local_method_name STRING
        [ local_model_pointer STRING ]
        )
      | local_method_pointer STRING
      [ local_search_probability REAL ]
      )
    |
    ( collaborative
      ( method_name_list STRINGLIST
        [ model_pointer_list STRING ]
        )
      | method_pointer_list STRINGLIST
      )
    [ iterator_servers INTEGER > 0 ]
    [ iterator_scheduling
      master
      | peer
      ]
    [ processors_per_iterator INTEGER > 0 ]
    )
  |
  ( multi_start
    ( method_name STRING
      [ model_pointer STRING ]
      )
    | method_pointer STRING
    [ random_starts INTEGER
      [ seed INTEGER ]
      ]
    [ starting_points REALLIST ]
    [ iterator_servers INTEGER > 0 ]
    [ iterator_scheduling
      master
      | peer
      ]
    [ processors_per_iterator INTEGER > 0 ]
```

```
    )
  |
  ( pareto_set
    ( method_name ALIAS opt_method_name STRING
      [ model_pointer ALIAS opt_model_pointer STRING ]
      )
    | method_pointer ALIAS opt_method_pointer STRING
    [ random_weight_sets INTEGER
      [ seed INTEGER ]
      ]
    [ weight_sets ALIAS multi_objective_weight_sets REALLIST ]
    [ iterator_servers INTEGER > 0 ]
    [ iterator_scheduling
      master
      | peer
      ]
    [ processors_per_iterator INTEGER > 0 ]
    )
  |
  ( surrogate_based_local
    method_pointer ALIAS approx_method_pointer STRING
    | method_name ALIAS approx_method_name STRING
    model_pointer ALIAS approx_model_pointer STRING
    [ soft_convergence_limit INTEGER ]
    [ truth_surrogate_bypass ]
    [ trust_region
      [ initial_size REAL ]
      [ minimum_size REAL ]
      [ contract_threshold REAL ]
      [ expand_threshold REAL ]
      [ contraction_factor REAL ]
      [ expansion_factor REAL ]
      ]
    [ approx_subproblem
      original_primary
      | single_objective
      | augmented_lagrangian_objective
      | lagrangian_objective
      original_constraints
      | linearized_constraints
      | no_constraints
      ]
    [ merit_function
      penalty_merit
      | adaptive_penalty_merit
      | lagrangian_merit
      | augmented_lagrangian_merit
      ]
    [ acceptance_logic
      tr_ratio
      | filter
      ]
    [ constraint_relax
      homotopy
      ]
    )
  |
  ( surrogate_based_global
    method_pointer ALIAS approx_method_pointer STRING
    | method_name ALIAS approx_method_name STRING
    model_pointer ALIAS approx_model_pointer STRING
    [ replace_points ]
```

```
  )
|
( dot_frcg
  [ linear_inequality_constraint_matrix REALLIST ]
  [ linear_inequality_lower_bounds REALLIST ]
  [ linear_inequality_upper_bounds REALLIST ]
  [ linear_inequality_scale_types STRINGLIST ]
  [ linear_inequality_scales REALLIST ]
  [ linear_equality_constraint_matrix REALLIST ]
  [ linear_equality_targets REALLIST ]
  [ linear_equality_scale_types STRINGLIST ]
  [ linear_equality_scales REALLIST ]
  [ model_pointer STRING ]
  )
| dot_mmfd
| dot_bfgs
| dot_slp
| dot_sqp
|
( dot
  frcg
  | mmfd
  | bfgs
  | slp
  | sqp
  [ linear_inequality_constraint_matrix REALLIST ]
  [ linear_inequality_lower_bounds REALLIST ]
  [ linear_inequality_upper_bounds REALLIST ]
  [ linear_inequality_scale_types STRINGLIST ]
  [ linear_inequality_scales REALLIST ]
  [ linear_equality_constraint_matrix REALLIST ]
  [ linear_equality_targets REALLIST ]
  [ linear_equality_scale_types STRINGLIST ]
  [ linear_equality_scales REALLIST ]
  [ model_pointer STRING ]
  )
|
( conmin_frcg
  [ linear_inequality_constraint_matrix REALLIST ]
  [ linear_inequality_lower_bounds REALLIST ]
  [ linear_inequality_upper_bounds REALLIST ]
  [ linear_inequality_scale_types STRINGLIST ]
  [ linear_inequality_scales REALLIST ]
  [ linear_equality_constraint_matrix REALLIST ]
  [ linear_equality_targets REALLIST ]
  [ linear_equality_scale_types STRINGLIST ]
  [ linear_equality_scales REALLIST ]
  [ model_pointer STRING ]
  )
| conmin_mfd
|
( conmin
  frcg
  | mfd
  [ linear_inequality_constraint_matrix REALLIST ]
  [ linear_inequality_lower_bounds REALLIST ]
  [ linear_inequality_upper_bounds REALLIST ]
  [ linear_inequality_scale_types STRINGLIST ]
  [ linear_inequality_scales REALLIST ]
  [ linear_equality_constraint_matrix REALLIST ]
  [ linear_equality_targets REALLIST ]
  [ linear_equality_scale_types STRINGLIST ]
```

```
    [ linear_equality_scales REALLIST ]
    [ model_pointer STRING ]
    )
  |
  ( dl_solver STRING
    [ linear_inequality_constraint_matrix REALLIST ]
    [ linear_inequality_lower_bounds REALLIST ]
    [ linear_inequality_upper_bounds REALLIST ]
    [ linear_inequality_scale_types STRINGLIST ]
    [ linear_inequality_scales REALLIST ]
    [ linear_equality_constraint_matrix REALLIST ]
    [ linear_equality_targets REALLIST ]
    [ linear_equality_scale_types STRINGLIST ]
    [ linear_equality_scales REALLIST ]
    [ model_pointer STRING ]
    )
  |
  ( npsol_sqp
    [ verify_level INTEGER ]
    [ function_precision REAL ]
    [ linesearch_tolerance REAL ]
    [ linear_inequality_constraint_matrix REALLIST ]
    [ linear_inequality_lower_bounds REALLIST ]
    [ linear_inequality_upper_bounds REALLIST ]
    [ linear_inequality_scale_types STRINGLIST ]
    [ linear_inequality_scales REALLIST ]
    [ linear_equality_constraint_matrix REALLIST ]
    [ linear_equality_targets REALLIST ]
    [ linear_equality_scale_types STRINGLIST ]
    [ linear_equality_scales REALLIST ]
    [ model_pointer STRING ]
    )
  | nlssol_sqp
  |
  ( stanford
    npsol
    | nlssol
    [ verify_level INTEGER ]
    [ function_precision REAL ]
    [ linesearch_tolerance REAL ]
    [ linear_inequality_constraint_matrix REALLIST ]
    [ linear_inequality_lower_bounds REALLIST ]
    [ linear_inequality_upper_bounds REALLIST ]
    [ linear_inequality_scale_types STRINGLIST ]
    [ linear_inequality_scales REALLIST ]
    [ linear_equality_constraint_matrix REALLIST ]
    [ linear_equality_targets REALLIST ]
    [ linear_equality_scale_types STRINGLIST ]
    [ linear_equality_scales REALLIST ]
    [ model_pointer STRING ]
    )
  |
  ( nlpql_sqp
    [ linear_inequality_constraint_matrix REALLIST ]
    [ linear_inequality_lower_bounds REALLIST ]
    [ linear_inequality_upper_bounds REALLIST ]
    [ linear_inequality_scale_types STRINGLIST ]
    [ linear_inequality_scales REALLIST ]
    [ linear_equality_constraint_matrix REALLIST ]
    [ linear_equality_targets REALLIST ]
    [ linear_equality_scale_types STRINGLIST ]
    [ linear_equality_scales REALLIST ]
```

```
    [ model_pointer STRING ]
    )
  |
  ( optpp_cg
    [ max_step REAL ]
    [ gradient_tolerance REAL ]
    [ linear_inequality_constraint_matrix REALLIST ]
    [ linear_inequality_lower_bounds REALLIST ]
    [ linear_inequality_upper_bounds REALLIST ]
    [ linear_inequality_scale_types STRINGLIST ]
    [ linear_inequality_scales REALLIST ]
    [ linear_equality_constraint_matrix REALLIST ]
    [ linear_equality_targets REALLIST ]
    [ linear_equality_scale_types STRINGLIST ]
    [ linear_equality_scales REALLIST ]
    [ model_pointer STRING ]
    )
  |
  ( optpp_q_newton
    | optpp_fd_newton
    | optpp_g_newton
    | optpp_newton
    [ search_method
      value_based_line_search
      | gradient_based_line_search
      | trust_region
      | tr_pds
      ]
    [ merit_function
      el_bakry
      | argaez_tapia
      | van_shanno
      ]
    [ steplength_to_boundary REAL ]
    [ centering_parameter REAL ]
    [ max_step REAL ]
    [ gradient_tolerance REAL ]
    [ linear_inequality_constraint_matrix REALLIST ]
    [ linear_inequality_lower_bounds REALLIST ]
    [ linear_inequality_upper_bounds REALLIST ]
    [ linear_inequality_scale_types STRINGLIST ]
    [ linear_inequality_scales REALLIST ]
    [ linear_equality_constraint_matrix REALLIST ]
    [ linear_equality_targets REALLIST ]
    [ linear_equality_scale_types STRINGLIST ]
    [ linear_equality_scales REALLIST ]
    [ model_pointer STRING ]
    )
  |
  ( optpp_pds
    [ search_scheme_size INTEGER ]
    [ linear_inequality_constraint_matrix REALLIST ]
    [ linear_inequality_lower_bounds REALLIST ]
    [ linear_inequality_upper_bounds REALLIST ]
    [ linear_inequality_scale_types STRINGLIST ]
    [ linear_inequality_scales REALLIST ]
    [ linear_equality_constraint_matrix REALLIST ]
    [ linear_equality_targets REALLIST ]
    [ linear_equality_scale_types STRINGLIST ]
    [ linear_equality_scales REALLIST ]
    [ model_pointer STRING ]
    )
```

```
|
( asynch_pattern_search ALIAS coliny_apps
  [ initial_delta REAL ]
  [ contraction_factor REAL ]
  [ threshold_delta REAL ]
  [ solution_target ALIAS solution_accuracy REAL ]
  [ synchronization
    blocking
    | nonblocking
    ]
  [ merit_function
    merit_max
    | merit_max_smooth
    | merit1
    | merit1_smooth
    | merit2
    | merit2_smooth
    | merit2_squared
    ]
  [ constraint_penalty REAL ]
  [ smoothing_factor REAL ]
  [ linear_inequality_constraint_matrix REALLIST ]
  [ linear_inequality_lower_bounds REALLIST ]
  [ linear_inequality_upper_bounds REALLIST ]
  [ linear_inequality_scale_types STRINGLIST ]
  [ linear_inequality_scales REALLIST ]
  [ linear_equality_constraint_matrix REALLIST ]
  [ linear_equality_targets REALLIST ]
  [ linear_equality_scale_types STRINGLIST ]
  [ linear_equality_scales REALLIST ]
  [ model_pointer STRING ]
  )
|
( mesh_adaptive_search
  [ function_precision REAL ]
  [ seed INTEGER > 0 ]
  [ history_file STRING ]
  [ display_format STRING ]
  [ variable_neighborhood_search REAL ]
  [ display_all_evaluations ]
  [ linear_inequality_constraint_matrix REALLIST ]
  [ linear_inequality_lower_bounds REALLIST ]
  [ linear_inequality_upper_bounds REALLIST ]
  [ linear_inequality_scale_types STRINGLIST ]
  [ linear_inequality_scales REALLIST ]
  [ linear_equality_constraint_matrix REALLIST ]
  [ linear_equality_targets REALLIST ]
  [ linear_equality_scale_types STRINGLIST ]
  [ linear_equality_scales REALLIST ]
  [ model_pointer STRING ]
  )
|
( moga
  [ fitness_type
    layer_rank
    | domination_count
    ]
  [ replacement_type
    elitist
    | roulette_wheel
    | unique_roulette_wheel
    |
```

```
    ( below_limit REAL
      [ shrinkage_fraction ALIAS shrinkage_percentage REAL ]
      )
    ]
  [ niching_type
    radial REALLIST
    | distance REALLIST
    |
    ( max_designs REALLIST
      [ num_designs INTEGER >= 2 ]
      )
    ]
  [ convergence_type
    metric_tracker
    [ percent_change REAL ]
    [ num_generations INTEGER >= 0 ]
    ]
  [ postprocessor_type
    orthogonal_distance REALLIST
    ]
  [ population_size INTEGER >= 0 ]
  [ log_file STRING ]
  [ print_each_pop ]
  [ initialization_type
    simple_random
    | unique_random
    | flat_file STRING
    ]
  [ crossover_type
    multi_point_binary INTEGER
    | multi_point_parameterized_binary INTEGER
    | multi_point_real INTEGER
    |
    ( shuffle_random
      [ num_parents INTEGER > 0 ]
      [ num_offspring INTEGER > 0 ]
      )
    [ crossover_rate REAL ]
    ]
  [ mutation_type
    bit_random
    | replace_uniform
    |
    ( offset_normal
      | offset_cauchy
      | offset_uniform
      [ mutation_scale REAL ]
      )
    [ mutation_rate REAL ]
    ]
  [ seed INTEGER > 0 ]
  [ linear_inequality_constraint_matrix REALLIST ]
  [ linear_inequality_lower_bounds REALLIST ]
  [ linear_inequality_upper_bounds REALLIST ]
  [ linear_inequality_scale_types STRINGLIST ]
  [ linear_inequality_scales REALLIST ]
  [ linear_equality_constraint_matrix REALLIST ]
  [ linear_equality_targets REALLIST ]
  [ linear_equality_scale_types STRINGLIST ]
  [ linear_equality_scales REALLIST ]
  [ model_pointer STRING ]
  )
```

```
|
( soga
  [ fitness_type
    merit_function
    [ constraint_penalty REAL ]
    ]
  [ replacement_type
    elitist
    | favor_feasible
    | roulette_wheel
    | unique_roulette_wheel
    ]
  [ convergence_type
    ( best_fitness_tracker
      [ percent_change REAL ]
      [ num_generations INTEGER >= 0 ]
      )
    |
    ( average_fitness_tracker
      [ percent_change REAL ]
      [ num_generations INTEGER >= 0 ]
      )
    ]
  [ population_size INTEGER >= 0 ]
  [ log_file STRING ]
  [ print_each_pop ]
  [ initialization_type
    simple_random
    | unique_random
    | flat_file STRING
    ]
  [ crossover_type
    multi_point_binary INTEGER
    | multi_point_parameterized_binary INTEGER
    | multi_point_real INTEGER
    |
    ( shuffle_random
      [ num_parents INTEGER > 0 ]
      [ num_offspring INTEGER > 0 ]
      )
    [ crossover_rate REAL ]
    ]
  [ mutation_type
    bit_random
    | replace_uniform
    |
    ( offset_normal
      | offset_cauchy
      | offset_uniform
      [ mutation_scale REAL ]
      )
    [ mutation_rate REAL ]
    ]
  [ seed INTEGER > 0 ]
  [ linear_inequality_constraint_matrix REALLIST ]
  [ linear_inequality_lower_bounds REALLIST ]
  [ linear_inequality_upper_bounds REALLIST ]
  [ linear_inequality_scale_types STRINGLIST ]
  [ linear_inequality_scales REALLIST ]
  [ linear_equality_constraint_matrix REALLIST ]
  [ linear_equality_targets REALLIST ]
  [ linear_equality_scale_types STRINGLIST ]
```

```
    [ linear_equality_scales REALLIST ]
    [ model_pointer STRING ]
    )
  |
  ( coliny_pattern_search
    [ constant_penalty ]
    [ no_expansion ]
    [ expand_after_success INTEGER ]
    [ pattern_basis
      coordinate
      | simplex
      ]
    [ stochastic ]
    [ total_pattern_size INTEGER ]
    [ exploratory_moves
      multi_step
      | adaptive_pattern
      | basic_pattern
      ]
    [ synchronization
      blocking
      | nonblocking
      ]
    [ contraction_factor REAL ]
    [ constraint_penalty REAL ]
    [ initial_delta REAL ]
    [ threshold_delta REAL ]
    [ solution_target ALIAS solution_accuracy REAL ]
    [ seed INTEGER > 0 ]
    [ show_misc_options ]
    [ misc_options STRINGLIST ]
    [ model_pointer STRING ]
    )
  |
  ( coliny_solis_wets
    [ contract_after_failure INTEGER ]
    [ no_expansion ]
    [ expand_after_success INTEGER ]
    [ constant_penalty ]
    [ contraction_factor REAL ]
    [ constraint_penalty REAL ]
    [ initial_delta REAL ]
    [ threshold_delta REAL ]
    [ solution_target ALIAS solution_accuracy REAL ]
    [ seed INTEGER > 0 ]
    [ show_misc_options ]
    [ misc_options STRINGLIST ]
    [ model_pointer STRING ]
    )
  |
  ( coliny_cobyla
    [ initial_delta REAL ]
    [ threshold_delta REAL ]
    [ solution_target ALIAS solution_accuracy REAL ]
    [ seed INTEGER > 0 ]
    [ show_misc_options ]
    [ misc_options STRINGLIST ]
    [ model_pointer STRING ]
    )
  |
  ( coliny_direct
    [ division
```

```
      major_dimension
      | all_dimensions
      ]
    [ global_balance_parameter REAL ]
    [ local_balance_parameter REAL ]
    [ max_boxsize_limit REAL ]
    [ min_boxsize_limit REAL ]
    [ constraint_penalty REAL ]
    [ solution_target ALIAS solution_accuracy REAL ]
    [ seed INTEGER > 0 ]
    [ show_misc_options ]
    [ misc_options STRINGLIST ]
    [ model_pointer STRING ]
    )
  |
  ( coliny_ea
    [ population_size INTEGER > 0 ]
    [ initialization_type
      simple_random
      | unique_random
      | flat_file STRING
      ]
    [ fitness_type
      linear_rank
      | merit_function
      ]
    [ replacement_type
      random INTEGER
      | chc INTEGER
      | elitist INTEGER
      [ new_solutions_generated INTEGER ]
      ]
    [ crossover_rate REAL ]
    [ crossover_type
      two_point
      | blend
      | uniform
      ]
    [ mutation_rate REAL ]
    [ mutation_type
      replace_uniform
      |
      ( offset_normal
        | offset_cauchy
        | offset_uniform
        [ mutation_scale REAL ]
        [ mutation_range INTEGER ]
        )
      [ non_adaptive ]
      ]
    [ constraint_penalty REAL ]
    [ solution_target ALIAS solution_accuracy REAL ]
    [ seed INTEGER > 0 ]
    [ show_misc_options ]
    [ misc_options STRINGLIST ]
    [ model_pointer STRING ]
    )
  |
  ( coliny_beta
    beta_solver_name STRING
    [ solution_target ALIAS solution_accuracy REAL ]
    [ seed INTEGER > 0 ]
```

```
    [ show_misc_options ]
    [ misc_options STRINGLIST ]
    [ model_pointer STRING ]
    )
  |
  ( nl2sol
    [ function_precision REAL ]
    [ absolute_conv_tol REAL ]
    [ x_conv_tol REAL ]
    [ singular_conv_tol REAL ]
    [ singular_radius REAL ]
    [ false_conv_tol REAL ]
    [ initial_trust_radius REAL ]
    [ covariance INTEGER ]
    [ regression_diagnostics ]
    [ model_pointer STRING ]
    )
  |
  ( nonlinear_cg
    [ misc_options STRINGLIST ]
    [ model_pointer STRING ]
    )
  |
  ( ncsu_direct
    [ solution_target ALIAS solution_accuracy REAL ]
    [ min_boxsize_limit REAL ]
    [ volume_boxsize_limit REAL ]
    [ model_pointer STRING ]
    )
  |
  ( genie_opt_darts
    | genie_direct
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( efficient_global
    [ gaussian_process ALIAS kriging
      surfpack
      | dakota
      ]
    [ use_derivatives ]
    [ import_points_file STRING
      [ annotated
      | freeform ]
      ]
    [ export_points_file STRING
      [ annotated
      | freeform ]
      ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( polynomial_chaos ALIAS nond_polynomial_chaos
    [ p_refinement
      uniform
      |
      ( dimension_adaptive
        sobol
        | decay
        | generalized
```

```
      )
    ]
  [ askey
  | wiener ]
  ( quadrature_order INTEGERLIST
    [ dimension_preference REALLIST ]
    [ nested
    | non_nested ]
    )
  |
  ( sparse_grid_level INTEGERLIST
    [ restricted
    | unrestricted ]
    [ dimension_preference REALLIST ]
    [ nested
    | non_nested ]
    )
  | cubature_integrand INTEGER
  |
  ( expansion_order INTEGERLIST
    [ dimension_preference REALLIST ]
    [ basis_type
      tensor_product
      | total_order
      | adapted
      ]
    ( collocation_points INTEGERLIST
      | collocation_ratio REAL
      [ ratio_order REAL ]
      [ ( least_squares
      [ svd
      | equality_constrained ]
      )
      |
      ( orthogonal_matching_pursuit ALIAS omp
    [ noise_tolerance REALLIST ]
    ]
      | basis_pursuit ALIAS bp
      |
      ( basis_pursuit_denoising ALIAS bpdn
    [ noise_tolerance REALLIST ]
    ]
      |
      ( least_angle_regression ALIAS lars
    [ noise_tolerance REALLIST ]
    ]
      |
      ( least_absolute_shrinkage ALIAS lasso
    [ noise_tolerance REALLIST ]
    [ l2_penalty REAL ]
    ]
      [ cross_validation ]
      [ use_derivatives ]
      [ tensor_grid ]
      [ reuse_points ALIAS reuse_samples ]
      )
    |
    ( expansion_samples INTEGERLIST
      [ reuse_points ALIAS reuse_samples ]
      [ incremental_lhs ]
      )
    [ import_points_file STRING
```

```
      [ annotated
      | freeform ]
      ]
    )
  |
  ( orthogonal_least_interpolation ALIAS least_interpolation ALIAS oli
    collocation_points INTEGERLIST
    [ cross_validation ]
    [ tensor_grid INTEGERLIST ]
    [ reuse_points ALIAS reuse_samples ]
    [ import_points_file STRING
      [ annotated
      | freeform ]
      ]
    )
  | import_expansion_file STRING
  [ variance_based_decomp
    [ interaction_order INTEGER > 0 ]
    [ drop_tolerance REAL ]
    ]
  [ diagonal_covariance
  | full_covariance ]
  [ normalized ]
  [ sample_type
    lhs
    | random
    ]
  [ probability_refinement ALIAS sample_refinement
    import
    | adapt_import
    | mm_adapt_import
    [ refinement_samples INTEGER ]
    ]
  [ export_points_file STRING
    [ annotated
    | freeform ]
    ]
  [ export_expansion_file STRING ]
  [ fixed_seed ]
  [ reliability_levels REALLIST
    [ num_reliability_levels INTEGERLIST ]
    ]
  [ response_levels REALLIST
    [ num_response_levels INTEGERLIST ]
    [ compute
      probabilities
      | reliabilities
      | gen_reliabilities
      [ system
    series
    | parallel
    ]
      ]
    ]
  [ distribution
    cumulative
    | complementary
    ]
  [ probability_levels REALLIST
    [ num_probability_levels INTEGERLIST ]
    ]
  [ gen_reliability_levels REALLIST
```

```
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ rng
      mt19937
      | rnum2
      ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( stoch_collocation ALIAS nond_stoch_collocation
    [ ( p_refinement
        uniform
        |
        ( dimension_adaptive
      sobol
      | generalized
      )
        )
    |
    ( h_refinement
      uniform
      |
      ( dimension_adaptive
        sobol
        | generalized
        )
      | local_adaptive
      ]
    [ piecewise
    | askey
    | wiener ]
    quadrature_order INTEGERLIST
    |
    ( sparse_grid_level INTEGERLIST
      [ restricted
      | unrestricted ]
      [ nodal
      | hierarchical ]
      )
    [ dimension_preference REALLIST ]
    [ use_derivatives ]
    [ nested
    | non_nested ]
    [ variance_based_decomp
      [ interaction_order INTEGER > 0 ]
      [ drop_tolerance REAL ]
      ]
    [ diagonal_covariance
    | full_covariance ]
    [ sample_type
      lhs
      | random
      ]
    [ probability_refinement ALIAS sample_refinement
      import
      | adapt_import
      | mm_adapt_import
      [ refinement_samples INTEGER ]
      ]
    [ export_points_file STRING
```

```
      [ annotated
      | freeform ]
      ]
    [ fixed_seed ]
    [ reliability_levels REALLIST
      [ num_reliability_levels INTEGERLIST ]
      ]
    [ response_levels REALLIST
      [ num_response_levels INTEGERLIST ]
      [ compute
        probabilities
        | reliabilities
        | gen_reliabilities
        [ system
      series
      | parallel
      ]
        ]
      ]
    [ distribution
      cumulative
      | complementary
      ]
    [ probability_levels REALLIST
      [ num_probability_levels INTEGERLIST ]
      ]
    [ gen_reliability_levels REALLIST
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ rng
      mt19937
      | rnum2
      ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( sampling ALIAS nond_sampling
    [ sample_type
      random
      | lhs
      |
      ( incremental_lhs
        | incremental_random
        previous_samples INTEGER
        )
      ]
    [ variance_based_decomp
      [ drop_tolerance REAL ]
      ]
    [ fixed_seed ]
    [ reliability_levels REALLIST
      [ num_reliability_levels INTEGERLIST ]
      ]
    [ response_levels REALLIST
      [ num_response_levels INTEGERLIST ]
      [ compute
        probabilities
        | reliabilities
        | gen_reliabilities
        [ system
```

```
      series
      | parallel
      ]
        ]
      ]
    [ distribution
      cumulative
      | complementary
      ]
    [ probability_levels REALLIST
      [ num_probability_levels INTEGERLIST ]
      ]
    [ gen_reliability_levels REALLIST
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ rng
      mt19937
      | rnum2
      ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( importance_sampling ALIAS nond_importance_sampling
    import
    | adapt_import
    | mm_adapt_import
    [ refinement_samples INTEGER ]
    [ response_levels REALLIST
      [ num_response_levels INTEGERLIST ]
      [ compute
        probabilities
        | gen_reliabilities
        [ system
      series
      | parallel
      ]
        ]
      ]
    [ distribution
      cumulative
      | complementary
      ]
    [ probability_levels REALLIST
      [ num_probability_levels INTEGERLIST ]
      ]
    [ gen_reliability_levels REALLIST
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ rng
      mt19937
      | rnum2
      ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( gpais ALIAS gaussian_process_adaptive_importance_sampling
    [ emulator_samples INTEGER ]
    [ import_points_file STRING
```

```
      [ annotated
      | freeform ]
      ]
    [ export_points_file STRING
      [ annotated
      | freeform ]
      ]
    [ response_levels REALLIST
      [ num_response_levels INTEGERLIST ]
      [ compute
        probabilities
        | gen_reliabilities
        [ system
      series
      | parallel
      ]
        ]
      ]
    [ distribution
      cumulative
      | complementary
      ]
    [ probability_levels REALLIST
      [ num_probability_levels INTEGERLIST ]
      ]
    [ gen_reliability_levels REALLIST
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ rng
      mt19937
      | rnum2
      ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( adaptive_sampling ALIAS nond_adaptive_sampling
    [ emulator_samples INTEGER ]
    [ fitness_metric
      predicted_variance
      | distance
      | gradient
      ]
    [ batch_selection
      naive
      | distance_penalty
      | topology
      | constant_liar
      ]
    [ batch_size INTEGER ]
    [ import_points_file STRING
      [ annotated
      | freeform ]
      ]
    [ export_points_file STRING
      [ annotated
      | freeform ]
      ]
    [ response_levels REALLIST
      [ num_response_levels INTEGERLIST ]
      [ compute
```

```
        probabilities
        | gen_reliabilities
        [ system
      series
      | parallel
      ]
        ]
      ]
    [ misc_options STRINGLIST ]
    [ distribution
      cumulative
      | complementary
      ]
    [ probability_levels REALLIST
      [ num_probability_levels INTEGERLIST ]
      ]
    [ gen_reliability_levels REALLIST
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ rng
      mt19937
      | rnum2
      ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( pof_darts ALIAS nond_pof_darts
    [ response_levels REALLIST
      [ num_response_levels INTEGERLIST ]
      [ compute
        probabilities
        | gen_reliabilities
        [ system
      series
      | parallel
      ]
        ]
      ]
    [ distribution
      cumulative
      | complementary
      ]
    [ probability_levels REALLIST
      [ num_probability_levels INTEGERLIST ]
      ]
    [ gen_reliability_levels REALLIST
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ rng
      mt19937
      | rnum2
      ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( efficient_subspace ALIAS nond_efficient_subspace
    [ emulator_samples INTEGER ]
    [ batch_size INTEGER ]
```

```
    [ distribution
      cumulative
      | complementary
      ]
    [ probability_levels REALLIST
      [ num_probability_levels INTEGERLIST ]
      ]
    [ gen_reliability_levels REALLIST
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ rng
      mt19937
      | rnum2
      ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( global_evidence ALIAS nond_global_evidence
    [ sbo
      | ego
      [ gaussian_process ALIAS kriging
        surfpack
        | dakota
        ]
      [ use_derivatives ]
      [ import_points_file STRING
        [ annotated
        | freeform ]
        ]
      [ export_points_file STRING
        [ annotated
        | freeform ]
        ]
      ]
    | ea
    | lhs ]
    [ response_levels REALLIST
      [ num_response_levels INTEGERLIST ]
      [ compute
        probabilities
        | gen_reliabilities
        [ system
      series
      | parallel
      ]
        ]
      ]
    [ distribution
      cumulative
      | complementary
      ]
    [ probability_levels REALLIST
      [ num_probability_levels INTEGERLIST ]
      ]
    [ gen_reliability_levels REALLIST
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ rng
      mt19937
      | rnum2
```

```
        ]
      [ samples INTEGER ]
      [ seed INTEGER > 0 ]
      [ model_pointer STRING ]
      )
    |
    ( global_interval_est ALIAS nond_global_interval_est
      [ sbo
        | ego
        [ gaussian_process ALIAS kriging
          surfpack
          | dakota
          ]
        [ use_derivatives ]
        [ import_points_file STRING
          [ annotated
          | freeform ]
          ]
        [ export_points_file STRING
          [ annotated
          | freeform ]
          ]
        ]
      | ea
      | lhs ]
      [ rng
        mt19937
        | rnum2
        ]
      [ samples INTEGER ]
      [ seed INTEGER > 0 ]
      [ model_pointer STRING ]
      )
    |
    ( bayes_calibration ALIAS nond_bayes_calibration
      ( queso
        [ emulator
          ( gaussian_process ALIAS kriging
        surfpack
        | dakota
        [ emulator_samples INTEGER ]
        [ import_points_file STRING
          [ annotated
          | freeform ]
          ]
        [ export_points_file STRING
          [ annotated
          | freeform ]
          ]
        )
          |
          ( pce
        [ sparse_grid_level INTEGERLIST ]
        )
          |
          ( sc
        [ sparse_grid_level INTEGERLIST ]
        )
          ]
        [ mcmc_type
          ( dram
        [ rejection
```

```
      standard
      | delayed
      ]
    [ metropolis
      hastings
      | adaptive
      ]
    )
      | multilevel
      ]
    [ rng
      mt19937
      | rnum2
      ]
    [ proposal_covariance_scale REALLIST ]
    )
  |
  ( gpmsa
    emulator_samples INTEGER
    [ import_points_file STRING
      [ annotated
      | freeform ]
      ]
    [ export_points_file STRING
      [ annotated
      | freeform ]
      ]
    [ mcmc_type
      ( dram
    [ rejection
      standard
      | delayed
      ]
    [ metropolis
      hastings
      | adaptive
      ]
    )
      | multilevel
      ]
    [ rng
      mt19937
      | rnum2
      ]
    [ proposal_covariance_scale REALLIST ]
    )
  |
  ( dream
    [ chains INTEGER >= 3 ]
    [ num_cr INTEGER >= 1 ]
    [ crossover_chain_pairs INTEGER >= 0 ]
    [ gr_threshold REAL > 0.0 ]
    [ jump_step INTEGER >= 0 ]
    [ emulator
      ( gaussian_process ALIAS kriging
    surfpack
    | dakota
    [ emulator_samples INTEGER ]
    [ import_points_file STRING
      [ annotated
      | freeform ]
      ]
```

```
      [ export_points_file STRING
        [ annotated
        | freeform ]
        ]
      )
        |
        ( pce
      [ sparse_grid_level INTEGERLIST ]
      )
        |
        ( sc
      [ sparse_grid_level INTEGERLIST ]
      )
        ]
      )
    [ use_derivatives ]
    [ likelihood_scale REAL ]
    [ calibrate_sigma ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( dace
    grid
    | random
    | oas
    | lhs
    | oa_lhs
    | box_behnken
    | central_composite
    [ main_effects ]
    [ quality_metrics ]
    [ variance_based_decomp
      [ drop_tolerance REAL ]
      ]
    [ fixed_seed ]
    [ symbols INTEGER ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( fsu_cvt
    [ latinize ]
    [ quality_metrics ]
    [ variance_based_decomp
      [ drop_tolerance REAL ]
      ]
    [ fixed_seed ]
    [ trial_type
      grid
      | halton
      | random
      ]
    [ num_trials INTEGER ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( psuade_moat
```

```
    [ partitions INTEGERLIST ]
    [ samples INTEGER ]
    [ seed INTEGER > 0 ]
    [ model_pointer STRING ]
    )
  |
  ( local_evidence ALIAS nond_local_evidence
    [ sqp
    | nip ]
    [ response_levels REALLIST
      [ num_response_levels INTEGERLIST ]
      [ compute
        probabilities
        | gen_reliabilities
        [ system
      series
      | parallel
      ]
        ]
      ]
    [ probability_levels REALLIST
      [ num_probability_levels INTEGERLIST ]
      ]
    [ gen_reliability_levels REALLIST
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ distribution
      cumulative
      | complementary
      ]
    [ model_pointer STRING ]
    )
  |
  ( local_interval_est ALIAS nond_local_interval_est
    [ sqp
    | nip ]
    [ model_pointer STRING ]
    )
  |
  ( local_reliability ALIAS nond_local_reliability
    [ mpp_search
      x_taylor_mean
      | u_taylor_mean
      | x_taylor_mpp
      | u_taylor_mpp
      | x_two_point
      | u_two_point
      | no_approx
      [ sqp
      | nip ]
      [ integration
        first_order
        | second_order
        [ probability_refinement ALIAS sample_refinement
      import
      | adapt_import
      | mm_adapt_import
      [ refinement_samples INTEGER ]
      [ seed INTEGER > 0 ]
      ]
        ]
      ]
```

```
    [ response_levels REALLIST
      [ num_response_levels INTEGERLIST ]
      [ compute
        probabilities
        | reliabilities
        | gen_reliabilities
        [ system
      series
      | parallel
      ]
        ]
      ]
    [ reliability_levels REALLIST
      [ num_reliability_levels INTEGERLIST ]
      ]
    [ distribution
      cumulative
      | complementary
      ]
    [ probability_levels REALLIST
      [ num_probability_levels INTEGERLIST ]
      ]
    [ gen_reliability_levels REALLIST
      [ num_gen_reliability_levels INTEGERLIST ]
      ]
    [ model_pointer STRING ]
    )
  |
  ( global_reliability ALIAS nond_global_reliability
    x_gaussian_process ALIAS x_kriging
    | u_gaussian_process ALIAS u_kriging
    [ surfpack
    | dakota ]
    [ import_points_file STRING
      [ annotated
      | freeform ]
      ]
    [ export_points_file STRING
      [ annotated
      | freeform ]
      ]
    [ use_derivatives ]
    [ seed INTEGER > 0 ]
    [ rng
      mt19937
      | rnum2
      ]
    [ response_levels REALLIST
      [ num_response_levels INTEGERLIST ]
      [ compute
        probabilities
        | gen_reliabilities
        [ system
      series
      | parallel
      ]
        ]
      ]
    [ distribution
      cumulative
      | complementary
      ]
```

```
      [ probability_levels REALLIST
        [ num_probability_levels INTEGERLIST ]
        ]
      [ gen_reliability_levels REALLIST
        [ num_gen_reliability_levels INTEGERLIST ]
        ]
      [ model_pointer STRING ]
      )
    |
    ( fsu_quasi_mc
      halton
      | hammersley
      [ latinize ]
      [ quality_metrics ]
      [ variance_based_decomp
        [ drop_tolerance REAL ]
        ]
      [ samples INTEGER ]
      [ fixed_sequence ]
      [ sequence_start INTEGERLIST ]
      [ sequence_leap INTEGERLIST ]
      [ prime_base INTEGERLIST ]
      [ model_pointer STRING ]
      )
    |
    ( vector_parameter_study
      final_point REALLIST
      | step_vector REALLIST
      num_steps INTEGER
      [ model_pointer STRING ]
      )
    |
    ( list_parameter_study
      list_of_points REALLIST
      |
      ( import_points_file STRING
        [ annotated
        | freeform ]
        )
      [ model_pointer STRING ]
      )
    |
    ( centered_parameter_study
      step_vector REALLIST
      steps_per_variable ALIAS deltas_per_variable INTEGERLIST
      [ model_pointer STRING ]
      )
    |
    ( multidim_parameter_study
      partitions INTEGERLIST
      [ model_pointer STRING ]
      )
    |
    ( richardson_extrap
      estimate_order
      | converge_order
      | converge_qoi
      [ refinement_rate REAL ]
      [ model_pointer STRING ]
      )

KEYWORD model
```

```
[ id_model STRING ]
[ variables_pointer STRING ]
[ responses_pointer STRING ]
[ hierarchical_tagging ]
( single
  [ interface_pointer STRING ]
  )
|
( surrogate
  [ id_surrogates INTEGERLIST ]
  ( global
    ( gaussian_process ALIAS kriging
      ( dakota
    [ point_selection ]
    [ trend
      constant
      | linear
      | reduced_quadratic
      ]
    )
      |
      ( surfpack
    [ trend
      constant
      | linear
      | reduced_quadratic
      | quadratic
      ]
    [ optimization_method STRING ]
    [ max_trials INTEGER > 0 ]
    [ nugget REAL > 0
    | find_nugget INTEGER ]
    [ correlation_lengths REALLIST ]
    [ export_model_file STRING ]
    )
      )
    |
    ( mars
      [ max_bases INTEGER ]
      [ interpolation
    linear
    | cubic
    ]
      [ export_model_file STRING ]
      )
    |
    ( moving_least_squares
      [ poly_order INTEGER ]
      [ weight_function INTEGER ]
      [ export_model_file STRING ]
      )
    |
    ( neural_network
      [ nodes INTEGER ]
      [ range REAL ]
      [ random_weight INTEGER ]
      [ export_model_file STRING ]
      )
    |
    ( radial_basis
      [ bases INTEGER ]
      [ max_pts INTEGER ]
```

```
      [ min_partition INTEGER ]
      [ max_subsets INTEGER ]
      [ export_model_file STRING ]
      )
    |
    ( polynomial
      linear
      | quadratic
      | cubic
      [ export_model_file STRING ]
      )
    [ total_points INTEGER
    | minimum_points
    | recommended_points ]
    [ dace_method_pointer STRING ]
    [ reuse_points ALIAS reuse_samples
      all
      | region
      | none
      ]
    [ import_points_file ALIAS samples_file STRING
      [ annotated
      | freeform ]
      ]
    [ export_points_file STRING
      [ annotated
      | freeform ]
      ]
    [ use_derivatives ]
    [ correction
      zeroth_order
      | first_order
      | second_order
      additive
      | multiplicative
      | combined
      ]
    [ metrics ALIAS diagnostics STRINGLIST
      [ cross_validate
    [ folds INTEGER
    | percent REAL ]
    ]
      [ press ]
      ]
    [ challenge_points_file STRING
      [ annotated
      | freeform ]
      ]
    )
  |
  ( multipoint
    tana
    actual_model_pointer STRING
    )
  |
  ( local
    taylor_series
    actual_model_pointer STRING
    )
  |
  ( hierarchical
    low_fidelity_model_pointer STRING
```

```
        high_fidelity_model_pointer STRING
        ( correction
          zeroth_order
          | first_order
          | second_order
          additive
          | multiplicative
          | combined
          )
        )
      )
    |
    ( nested
      [ optional_interface_pointer STRING
        [ optional_interface_responses_pointer STRING ]
        ]
      ( sub_method_pointer STRING
        [ iterator_servers INTEGER > 0 ]
        [ iterator_scheduling
          master
          | peer
          ]
        [ processors_per_iterator INTEGER > 0 ]
        [ primary_variable_mapping STRINGLIST ]
        [ secondary_variable_mapping STRINGLIST ]
        [ primary_response_mapping REALLIST ]
        [ secondary_response_mapping REALLIST ]
        )
      )

KEYWORD12 variables
    [ id_variables STRING ]
    [ active
      all
      | design
      | uncertain
      | aleatory
      | epistemic
      | state
      ]
    [ mixed
    | relaxed ]
    [ continuous_design INTEGER > 0
      [ initial_point ALIAS cdv_initial_point REALLIST ]
      [ lower_bounds ALIAS cdv_lower_bounds REALLIST ]
      [ upper_bounds ALIAS cdv_upper_bounds REALLIST ]
      [ scale_types ALIAS cdv_scale_types STRINGLIST ]
      [ scales ALIAS cdv_scales REALLIST ]
      [ descriptors ALIAS cdv_descriptors STRINGLIST ]
      ]
    [ discrete_design_range INTEGER > 0
      [ initial_point ALIAS ddv_initial_point INTEGERLIST ]
      [ lower_bounds ALIAS ddv_lower_bounds INTEGERLIST ]
      [ upper_bounds ALIAS ddv_upper_bounds INTEGERLIST ]
      [ descriptors ALIAS ddv_descriptors STRINGLIST ]
      ]
    [ discrete_design_set_integer INTEGER > 0
      [ initial_point INTEGERLIST ]
      [ num_set_values INTEGERLIST ]
      set_values INTEGERLIST
      [ descriptors STRINGLIST ]
      ]
```

```
[ discrete_design_set_real INTEGER > 0
  [ initial_point REALLIST ]
  [ num_set_values INTEGERLIST ]
  set_values REALLIST
  [ descriptors STRINGLIST ]
  ]
[ normal_uncertain INTEGER > 0
  means ALIAS nuv_means REALLIST
  std_deviations ALIAS nuv_std_deviations REALLIST
  [ lower_bounds ALIAS nuv_lower_bounds REALLIST ]
  [ upper_bounds ALIAS nuv_upper_bounds REALLIST ]
  [ initial_point REALLIST ]
  [ descriptors ALIAS nuv_descriptors STRINGLIST ]
  ]
[ lognormal_uncertain INTEGER > 0
  ( lambdas ALIAS lnuv_lambdas REALLIST
    zetas ALIAS lnuv_zetas REALLIST
    )
  |
  ( means ALIAS lnuv_means REALLIST
    std_deviations ALIAS lnuv_std_deviations REALLIST
    | error_factors ALIAS lnuv_error_factors REALLIST
    )
  [ lower_bounds ALIAS lnuv_lower_bounds REALLIST ]
  [ upper_bounds ALIAS lnuv_upper_bounds REALLIST ]
  [ initial_point REALLIST ]
  [ descriptors ALIAS lnuv_descriptors STRINGLIST ]
  ]
[ uniform_uncertain INTEGER > 0
  lower_bounds ALIAS uuv_lower_bounds REALLIST
  upper_bounds ALIAS uuv_upper_bounds REALLIST
  [ initial_point REALLIST ]
  [ descriptors ALIAS uuv_descriptors STRINGLIST ]
  ]
[ loguniform_uncertain INTEGER > 0
  lower_bounds ALIAS luuv_lower_bounds REALLIST
  upper_bounds ALIAS luuv_upper_bounds REALLIST
  [ initial_point REALLIST ]
  [ descriptors ALIAS luuv_descriptors STRINGLIST ]
  ]
[ triangular_uncertain INTEGER > 0
  modes ALIAS tuv_modes REALLIST
  lower_bounds ALIAS tuv_lower_bounds REALLIST
  upper_bounds ALIAS tuv_upper_bounds REALLIST
  [ initial_point REALLIST ]
  [ descriptors ALIAS tuv_descriptors STRINGLIST ]
  ]
[ exponential_uncertain INTEGER > 0
  betas ALIAS euv_betas REALLIST
  [ initial_point REALLIST ]
  [ descriptors ALIAS euv_descriptors STRINGLIST ]
  ]
[ beta_uncertain INTEGER > 0
  alphas ALIAS buv_alphas REALLIST
  betas ALIAS buv_betas REALLIST
  lower_bounds ALIAS buv_lower_bounds REALLIST
  upper_bounds ALIAS buv_upper_bounds REALLIST
  [ initial_point REALLIST ]
  [ descriptors ALIAS buv_descriptors STRINGLIST ]
  ]
[ gamma_uncertain INTEGER > 0
  alphas ALIAS gauv_alphas REALLIST
```

```
    betas ALIAS gauv_betas REALLIST
    [ initial_point REALLIST ]
    [ descriptors ALIAS gauv_descriptors STRINGLIST ]
    ]
  [ gumbel_uncertain INTEGER > 0
    alphas ALIAS guuv_alphas REALLIST
    betas ALIAS guuv_betas REALLIST
    [ initial_point REALLIST ]
    [ descriptors ALIAS guuv_descriptors STRINGLIST ]
    ]
  [ frechet_uncertain INTEGER > 0
    alphas ALIAS fuv_alphas REALLIST
    betas ALIAS fuv_betas REALLIST
    [ initial_point REALLIST ]
    [ descriptors ALIAS fuv_descriptors STRINGLIST ]
    ]
  [ weibull_uncertain INTEGER > 0
    alphas ALIAS wuv_alphas REALLIST
    betas ALIAS wuv_betas REALLIST
    [ initial_point REALLIST ]
    [ descriptors ALIAS wuv_descriptors STRINGLIST ]
    ]
  [ histogram_bin_uncertain INTEGER > 0
    [ num_pairs ALIAS huv_num_bin_pairs INTEGERLIST ]
    abscissas ALIAS huv_bin_abscissas REALLIST
    ordinates ALIAS huv_bin_ordinates REALLIST
    | counts ALIAS huv_bin_counts REALLIST
    [ initial_point REALLIST ]
    [ descriptors ALIAS huv_bin_descriptors STRINGLIST ]
    ]
  [ poisson_uncertain INTEGER > 0
    lambdas REALLIST
    [ initial_point INTEGERLIST ]
    [ descriptors STRINGLIST ]
    ]
  [ binomial_uncertain INTEGER > 0
    probability_per_trial ALIAS prob_per_trial REALLIST
    num_trials INTEGERLIST
    [ initial_point INTEGERLIST ]
    [ descriptors STRINGLIST ]
    ]
  [ negative_binomial_uncertain INTEGER > 0
    probability_per_trial ALIAS prob_per_trial REALLIST
    num_trials INTEGERLIST
    [ initial_point INTEGERLIST ]
    [ descriptors STRINGLIST ]
    ]
  [ geometric_uncertain INTEGER > 0
    probability_per_trial ALIAS prob_per_trial REALLIST
    [ initial_point INTEGERLIST ]
    [ descriptors STRINGLIST ]
    ]
  [ hypergeometric_uncertain INTEGER > 0
    total_population INTEGERLIST
    selected_population INTEGERLIST
    num_drawn INTEGERLIST
    [ initial_point INTEGERLIST ]
    [ descriptors STRINGLIST ]
    ]
  [ histogram_point_uncertain INTEGER > 0
    [ num_pairs ALIAS huv_num_point_pairs INTEGERLIST ]
    abscissas ALIAS huv_point_abscissas REALLIST
```

```
      counts ALIAS huv_point_counts REALLIST
      [ initial_point REALLIST ]
      [ descriptors ALIAS huv_point_descriptors STRINGLIST ]
      ]
    [ uncertain_correlation_matrix REALLIST ]
    [ continuous_interval_uncertain ALIAS interval_uncertain INTEGER > 0
      [ num_intervals ALIAS iuv_num_intervals INTEGERLIST ]
      [ interval_probabilities ALIAS interval_probs ALIAS iuv_interval_probs REALLIST ]
      lower_bounds REALLIST
      upper_bounds REALLIST
      [ initial_point REALLIST ]
      [ descriptors ALIAS iuv_descriptors STRINGLIST ]
      ]
    [ discrete_interval_uncertain ALIAS discrete_uncertain_range INTEGER > 0
      [ num_intervals INTEGERLIST ]
      [ interval_probabilities ALIAS interval_probs ALIAS range_probabilities ALIAS range_probs REALLIST ]
      lower_bounds INTEGERLIST
      upper_bounds INTEGERLIST
      [ initial_point INTEGERLIST ]
      [ descriptors STRINGLIST ]
      ]
    [ discrete_uncertain_set_integer INTEGER > 0
      [ num_set_values INTEGERLIST ]
      set_values INTEGERLIST
      [ set_probabilities ALIAS set_probs REALLIST ]
      [ initial_point INTEGERLIST ]
      [ descriptors STRINGLIST ]
      ]
    [ discrete_uncertain_set_real INTEGER > 0
      [ num_set_values INTEGERLIST ]
      set_values REALLIST
      [ set_probabilities ALIAS set_probs REALLIST ]
      [ initial_point REALLIST ]
      [ descriptors STRINGLIST ]
      ]
    [ continuous_state INTEGER > 0
      [ initial_state ALIAS csv_initial_state REALLIST ]
      [ lower_bounds ALIAS csv_lower_bounds REALLIST ]
      [ upper_bounds ALIAS csv_upper_bounds REALLIST ]
      [ descriptors ALIAS csv_descriptors STRINGLIST ]
      ]
    [ discrete_state_range INTEGER > 0
      [ initial_state ALIAS dsv_initial_state INTEGERLIST ]
      [ lower_bounds ALIAS dsv_lower_bounds INTEGERLIST ]
      [ upper_bounds ALIAS dsv_upper_bounds INTEGERLIST ]
      [ descriptors ALIAS dsv_descriptors STRINGLIST ]
      ]
    [ discrete_state_set_integer INTEGER > 0
      [ initial_state INTEGERLIST ]
      [ num_set_values INTEGERLIST ]
      set_values INTEGERLIST
      [ descriptors STRINGLIST ]
      ]
    [ discrete_state_set_real INTEGER > 0
      [ initial_state REALLIST ]
      [ num_set_values INTEGERLIST ]
      set_values REALLIST
      [ descriptors STRINGLIST ]
      ]

KEYWORD12 interface
    [ id_interface STRING ]
```

```
[ algebraic_mappings STRING ]
[ analysis_drivers STRINGLIST
  [ analysis_components STRINGLIST ]
  [ input_filter STRING ]
  [ output_filter STRING ]
  ( system
    | fork
    [ parameters_file STRING ]
    [ results_file STRING ]
    [ allow_existing_results ]
    [ verbatim ]
    [ aprepro ALIAS dprepro ]
    [ file_tag ]
    [ file_save ]
    [ work_directory
      [ named STRING ]
      [ directory_tag ALIAS dir_tag ]
      [ directory_save ALIAS dir_save ]
      [ template_directory STRING
    | template_files STRINGLIST
    [ copy ]
    [ replace ]
    ]
      ]
    )
  |
  ( direct
    [ processors_per_analysis INTEGER > 0 ]
    )
  | matlab
  |
  ( python
    [ numpy ]
    )
  | scilab
  | grid
  [ failure_capture
    abort
    | retry INTEGER
    | recover REALLIST
    | continuation
    ]
  [ deactivate
    [ active_set_vector ]
    [ evaluation_cache ]
    [ strict_cache_equality
      [ cache_tolerance REAL ]
      ]
    [ restart_file ]
    ]
  ]
[ asynchronous
  [ evaluation_concurrency INTEGER > 0 ]
  [ local_evaluation_scheduling
    dynamic
    | static
    ]
  [ analysis_concurrency INTEGER > 0 ]
  ]
[ evaluation_servers INTEGER > 0 ]
[ evaluation_scheduling
  master
```

```
      |
      ( peer
        dynamic
        | static
        )
      ]
    [ processors_per_evaluation INTEGER > 0 ]
    [ analysis_servers INTEGER > 0 ]
    [ analysis_scheduling
      master
      | peer
      ]

KEYWORD12 responses
    [ id_responses STRING ]
    [ descriptors ALIAS response_descriptors STRINGLIST ]
    ( objective_functions ALIAS num_objective_functions INTEGER >= 0
      [ sense STRINGLIST ]
      [ primary_scale_types ALIAS objective_function_scale_types STRINGLIST ]
      [ primary_scales ALIAS objective_function_scales REALLIST ]
      [ weights ALIAS multi_objective_weights REALLIST ]
      [ nonlinear_inequality_constraints ALIAS num_nonlinear_inequality_constraints INTEGER >= 0
        [ lower_bounds ALIAS nonlinear_inequality_lower_bounds REALLIST ]
        [ upper_bounds ALIAS nonlinear_inequality_upper_bounds REALLIST ]
        [ scale_types ALIAS nonlinear_inequality_scale_types STRINGLIST ]
        [ scales ALIAS nonlinear_inequality_scales REALLIST ]
        ]
      [ nonlinear_equality_constraints ALIAS num_nonlinear_equality_constraints INTEGER >= 0
        [ targets ALIAS nonlinear_equality_targets REALLIST ]
        [ scale_types ALIAS nonlinear_equality_scale_types STRINGLIST ]
        [ scales ALIAS nonlinear_equality_scales REALLIST ]
        ]
      )
    |
    ( calibration_terms ALIAS least_squares_terms ALIAS num_least_squares_terms INTEGER >= 0
      [ primary_scale_types ALIAS calibration_term_scale_types ALIAS least_squares_term_scale_types STRINGLIST ]
      [ primary_scales ALIAS calibration_term_scales ALIAS least_squares_term_scales REALLIST ]
      [ weights ALIAS calibration_weights ALIAS least_squares_weights REALLIST ]
      [ calibration_data_file ALIAS least_squares_data_file STRING
        [ num_experiments INTEGER >= 0 ]
        [ num_replicates INTEGERLIST ]
        [ annotated
        | freeform ]
        [ num_config_variables INTEGER >= 0 ]
        [ num_std_deviations INTEGER >= 0 ]
        ]
      [ nonlinear_inequality_constraints ALIAS num_nonlinear_inequality_constraints INTEGER >= 0
        [ lower_bounds ALIAS nonlinear_inequality_lower_bounds REALLIST ]
        [ upper_bounds ALIAS nonlinear_inequality_upper_bounds REALLIST ]
        [ scale_types ALIAS nonlinear_inequality_scale_types STRINGLIST ]
        [ scales ALIAS nonlinear_inequality_scales REALLIST ]
        ]
      [ nonlinear_equality_constraints ALIAS num_nonlinear_equality_constraints INTEGER >= 0
        [ targets ALIAS nonlinear_equality_targets REALLIST ]
        [ scale_types ALIAS nonlinear_equality_scale_types STRINGLIST ]
        [ scales ALIAS nonlinear_equality_scales REALLIST ]
        ]
      )
    | response_functions ALIAS num_response_functions INTEGER >= 0
    no_gradients
    | analytic_gradients
    |
```

```
( mixed_gradients
  id_numerical_gradients INTEGERLIST
  id_analytic_gradients INTEGERLIST
  [ method_source ]
  [ ( dakota
      [ ignore_bounds ]
      [ relative
      | absolute
      | bounds ]
      )
  | vendor ]
  [ interval_type ]
  [ forward
  | central ]
  [ fd_step_size ALIAS fd_gradient_step_size REALLIST ]
  )
|
( numerical_gradients
  [ method_source ]
  [ ( dakota
      [ ignore_bounds ]
      [ relative
      | absolute
      | bounds ]
      )
  | vendor ]
  [ interval_type ]
  [ forward
  | central ]
  [ fd_step_size ALIAS fd_gradient_step_size REALLIST ]
  )
no_hessians
|
( numerical_hessians
  [ fd_step_size ALIAS fd_hessian_step_size REALLIST ]
  [ relative
  | absolute
  | bounds ]
  [ forward
  | central ]
  )
|
( quasi_hessians
  ( bfgs
    [ damped ]
    )
  | sr1
  )
| analytic_hessians
|
( mixed_hessians
  [ id_numerical_hessians INTEGERLIST
    [ fd_step_size ALIAS fd_hessian_step_size REALLIST ]
    ]
  [ relative
  | absolute
  | bounds ]
  [ forward
  | central ]
  [ id_quasi_hessians INTEGERLIST
    ( bfgs
      [ damped ]
```

```
    )
  | sr1
  ]
[ id_analytic_hessians INTEGERLIST ]
)
```

Chapter 4

Topics Area

This page introduces the user to the topics used to organize keywords.

- admin
- dakota_IO
- dakota_concepts
- models
- variables_and_responses
- interface
- methods
- advanced_topics
- packages

4.1 admin

Description

This is only for management while ref man is under construction

Related Topics

- empty
- problem
- not_yet_reviewed

Related Keywords

4.1.1 empty

Description

This topic tracks the keywords which do not have content in the reference manual

Related Topics

Related Keywords

- misc_options :
- mcmc_type :
- dram :
- multilevel :
- non_adaptive :
- genie_direct :
- genie_opt_darts :
- global_model_pointer :
- local_model_pointer :
- misc_options :
- optional_interface_pointer :
- optional_interface_responses_pointer :
- primary_response_mapping :
- primary_variable_mapping :
- secondary_response_mapping :
- secondary_variable_mapping :
- point_selection :
- bases :
- high_fidelity_model_pointer :
- low_fidelity_model_pointer :
- set_probabilities :
- set_probabilities :

4.1.2 problem

Description

empty

Related Topics

Related Keywords

- interface : Specifies how function evaluations will be performed in order to map the variables into the responses.
- analysis_drivers : Define how Dakota should run a function evaluation
- analysis_components : Provide additional identifiers to analysis drivers.
- deactivate : Deactivate Dakota features to simplify interface development, increase execution speed, or reduce memory and disk requirements
- failure_capture : Determine how Dakota responds to analysis driver failure
- continuation : Cause Dakota to step toward the failed "target" simulation from a nearby successful "source"
- verbatim : Specify the command Dakota uses to launch analysis driver(s) and filters
- verbatim : Specify the command Dakota uses to launch analysis driver(s) and filters
- analysis_servers : Specify the number of analysis servers when Dakota is run in parallel
- asynchronous : Specify analysis driver concurrency, when Dakota is run in serial
- analysis_concurrency : Limit the number of analysis drivers within an evaluation that Dakota will schedule
- evaluation_concurrency : Determine how many concurrent evaluations Dakota will schedule
- evaluation_servers : Specify the number of evaluation servers when Dakota is run in parallel
- constant_liar : Use information from the existing surrogate model to predict what the surrogate upgrade will be with new points.
- topology : In this selection strategy, we use information about the topology of the space from the Morse--Smale complex to identify next points to select.
- system : ?????
- parallel : ?????
- series : ?????
- solution_target : Stopping criteria based on objective function value
- proposal_covariance_scale : Provide values for the diagonals of the proposal covariance matrix.
- proposal_covariance_scale : Provide values for the diagonals of the proposal covariance matrix.
- coliny_beta : What does this do????????????????
- solution_target : Stopping criteria based on objective function value
- solution_target : Stopping criteria based on objective function value
- solution_target : Stopping criteria based on objective function value
- initialization_type : Specify how to initialize the population

- replacement_type : Select a replacement type for JEGA methods
- solution_target : Stopping criteria based on objective function value
- solution_target : Stopping criteria based on objective function value
- solution_target : Stopping criteria based on objective function value
- dl_solver : There is ZERO mention of this in the reference manual??
- efficient_subspace : This is not mentioned in the reference manual
- global_evidence : Evidence theory with evidence measures computed with global optimization methods
- system : ?????
- parallel : ?????
- series : ?????
- global_reliability : Global reliability methods
- system : ?????
- parallel : ?????
- series : ?????
- system : ?????
- parallel : ?????
- series : ?????
- collaborative : Multiple methods run concurrently and share information
- embedded : A subordinate local method provides periodic refinements to a top-level global method
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- system : ?????
- parallel : ?????
- series : ?????
- nip : Use a nonlinear interior point method for optimization
- system : ?????
- parallel : ?????
- series : ?????
- sqp : Use a SQP method for optimization
- nip : Use a nonlinear interior point method for optimization
- sqp : Use a SQP method for optimization

- probability_refinement : Allow refinement of probability and generalized reliability results using importance sampling
- nip : Use a nonlinear interior point method for optimization
- sqp : Use a SQP method for optimization
- system : ?????
- parallel : ?????
- series : ?????
- metric_tracker : Track changes in the non-dominated frontier
- initialization_type : Specify how to initialize the population
- postprocessor_type : Post process the final solution from `moga`
- replacement_type : Select a replacement type for JEGA methods
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- method_name : Name of a sub-method to apply to surrogate model
- method_pointer : Pointer to sub-method to apply to surrogate
- verify_level : Verify the quality of analytic gradients
- nonlinear_cg : This is not mentioned in the reference manual
- verify_level : Verify the quality of analytic gradients
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- method_name : Name of a sub-method to apply to surrogate model
- method_pointer : Pointer to sub-method to apply to surrogate
- system : ?????
- parallel : ?????
- series : ?????
- reuse_points : ??
- reuse_points : ??
- reuse_points : ??
- reuse_points : ??
- probability_refinement : Allow refinement of probability and generalized reliability results using importance sampling
- system : ?????
- parallel : ?????

- series : ?????
- psuade_moat : Morris One-at-a-Time
- system : ?????
- parallel : ?????
- series : ?????
- initialization_type : Specify how to initialize the population
- replacement_type : Select a replacement type for JEGA methods
- favor_feasible : Prioritize feasible designs
- verify_level : Verify the quality of analytic gradients
- probability_refinement : Allow refinement of probability and generalized reliability results using importance sampling
- system : ?????
- parallel : ?????
- series : ?????
- method_name : Name of a sub-method to apply to surrogate model
- method_pointer : Pointer to sub-method to apply to surrogate
- method_name : Name of a sub-method to apply to surrogate model
- method_pointer : Pointer to sub-method to apply to surrogate
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- export_model_file : ????
- export_model_file : ????
- metrics : Compute metrics of surrogate quality
- cross_validate : k-fold cross validation
- export_model_file : ????
- export_model_file : ????
- nodes : ????
- random_weight : ????
- range : ????
- export_model_file : ????
- export_model_file : ????
- max_pts : ????

- max_subsets : ????
- min_partition : ????
- reuse_points : Specify how to use existing training data
- calibration_data_file : Supply calibration data in a text file
- counts : ???
- num_pairs : ???
- ordinates : ??????????

4.1.3 not_yet_reviewed

Description

Not yet reviewed.

Here's a test of the anchors and links
eq-correct_val_add
test caption
eq-taylor_a test caption 2 without quotes
eq-taylor_b
eq-exact_B
eq-exact_A
eq-taylor1
eq-taylor2

Related Topics

Related Keywords

- method : Specifies the name and behavior of a method
- batch_size : The number of points to add in each batch.
- system : ?????
- parallel : ?????
- series : ?????
- rnum2 : Random Number Generator 2
- synchronization : Select how Dakota schedules function evaluations
- blocking : The `blocking` option causes Dakota to run the algorithm in synchronous mode
- nonblocking : The `nonblocking` option causes Dakota to run the algorithm in asynchronous mode
- delayed : Select delayed rejection
- rnum2 : Random Number Generator 2
- likelihood_scale : Scale the **log-likelihood** function

- delayed : Select delayed rejection
- rnum2 : Random Number Generator 2
- coliny_beta : What does this do??????????????
- simple_random : Initilization option for JEGA methods
- replacement_type : Select a replacement type for JEGA methods
- synchronization : Select how Dakota schedules function evaluations
- blocking : The `blocking` option causes Dakota to run the algorithm in synchronous mode
- nonblocking : The `nonblocking` option causes Dakota to run the algorithm in asynchronous mode
- dace : Design and Analysis of Computer Experiments
- box_behnken : Box-Behnken Design
- central_composite : Central Composite Design
- grid : Grid Sampling
- main_effects : ANOVA
- oa_lhs : Orthogonal Array Latin Hypercube Sampling
- oas : Orthogonal Array Sampling
- quality_metrics : Calculate metrics to assess the quality of the samples
- symbols : Number of replications in the sample set
- dl_solver : There is ZERO mention of this in the reference manual??
- efficient_subspace : This is not mentioned in the reference manual
- batch_size : The number of points to add in each batch.
- rnum2 : Random Number Generator 2
- fsu_cvt : Design of Computer Experiments - Centroidal Voronoi Tessellation
- latinize : Adjust samples to improve the discrepancy of the marginal distributions
- num_trials : The number of secondary sample points generated to adjust the location of the primary sample points
- quality_metrics : Calculate metrics to assess the quality of the samples
- trial_type : Specify how the trial samples are generated
- grid : Samples on a regular grid
- halton : Generate samples from a Halton sequence
- fsu_quasi_mc : Design of Computer Experiments - Quasi-Monte Carlo sampling
- fixed_sequence : Reuse the same sequence and samples for multiple sampling sets

- halton : Generate samples from a Halton sequence
- hammersley : Use Hammersley sequences
- latinize : Adjust samples to improve the discrepancy of the marginal distributions
- prime_base : The prime numbers used to generate the sequence
- quality_metrics : Calculate metrics to assess the quality of the samples
- sequence_leap : Specify how often the sequence is sampled
- sequence_start : Choose where to start sampling the sequence
- system : ?????
- parallel : ?????
- series : ?????
- rnum2 : Random Number Generator 2
- rnum2 : Random Number Generator 2
- system : ?????
- parallel : ?????
- series : ?????
- rnum2 : Random Number Generator 2
- system : ?????
- parallel : ?????
- series : ?????
- rnum2 : Random Number Generator 2
- method_name_list :
- model_pointer_list :
- method_pointer_list :
- global_method_name : Specify the global method
- local_method_name : Specify the local method
- method_name_list :
- model_pointer_list :
- method_pointer_list :
- system : ?????
- parallel : ?????

- series : ?????
- rnum2 : Random Number Generator 2
- nip : Use a nonlinear interior point method for optimization
- system : ?????
- parallel : ?????
- series : ?????
- sqp : Use a SQP method for optimization
- nip : Use a nonlinear interior point method for optimization
- sqp : Use a SQP method for optimization
- nip : Use a nonlinear interior point method for optimization
- sqp : Use a SQP method for optimization
- system : ?????
- parallel : ?????
- series : ?????
- simple_random : Initilization option for JEGA methods
- log_file : Specify the name of a log file
- replacement_type : Select a replacement type for JEGA methods
- verify_level : Verify the quality of analytic gradients
- nonlinear_cg : This is not mentioned in the reference manual
- verify_level : Verify the quality of analytic gradients
- system : ?????
- parallel : ?????
- series : ?????
- rnum2 : Random Number Generator 2
- system : ?????
- parallel : ?????
- series : ?????
- rnum2 : Random Number Generator 2
- psuade_moat : Morris One-at-a-Time
- partitions : Number of partitions of each variable

- refinement_rate : Rate at which the state variables are refined
- system : ?????
- parallel : ?????
- series : ?????
- rnum2 : Random Number Generator 2
- scaling : Turn on scaling for variables, responses, and constraints
- simple_random : Initilization option for JEGA methods
- log_file : Specify the name of a log file
- replacement_type : Select a replacement type for JEGA methods
- stanford : Select methods from the Stanford package
- nlssol : Duplicate of `method-nlssol_sqp`
- npsol : Duplicate of `method-npsol_sqp`
- verify_level : Verify the quality of analytic gradients
- system : ?????
- parallel : ?????
- series : ?????
- rnum2 : Random Number Generator 2
- acceptance_logic : Set criteria for trusted surrogate
- model : Specifies how variables are mapped into a set of responses
- nested : A model whose responses are computed through the use of a sub-iterator
- sub_method_pointer : The `sub_method_pointer` specifies the method block for the sub-iterator
- interval_probabilities : Assign probability mass to each interval
- num_intervals : Specify the number of intervals for each variable
- num_set_values : The size of each variables' set of permissible values
- num_set_values : The size of each variables' set of permissible values
- interval_probabilities : Assign probability mass to each interval
- num_intervals : Specify the number of intervals for each variable
- num_set_values : The size of each variables' set of permissible values
- num_set_values : The size of each variables' set of permissible values
- num_set_values : The size of each variables' set of permissible values
- num_set_values : The size of each variables' set of permissible values

4.2 dakota_IO

Description

Keywords and Concepts relating inputs to Dakota and outputs from Dakota

Related Topics

- dakota_inputs
- dakota_output
- file_formats

Related Keywords

4.2.1 dakota_inputs

Description

empty

Related Topics

- block
- data_import_capabilities

Related Keywords

4.2.2 block

Description

A block is the highest level of keyword organization in Dakota. There are currently 6 blocks in the Dakota input spec:

Related Topics

- block_identifier
- block_pointer

Related Keywords

- environment : Top-level settings for Dakota execution
- interface : Specifies how function evaluations will be performed in order to map the variables into the responses.
- method : Specifies the name and behavior of a method
- model : Specifies how variables are mapped into a set of responses

- responses : Description of the model output data returned to Dakota upon evaluation of an interface.
- variables : Specifies the parameter set to be iterated by a particular method.

block_identifier

Description

empty

Related Topics

Related Keywords

- id_interface : Give the block a name, in case of multiple interface blocks
- id_method : Give the block a name, in case of multiple method blocks
- id_model : Give the model block an identifying name, in case of multiple model blocks
- id_responses : Name the response block, when there are multiple responses blocks

block_pointer

Description

See block_pointer for details about pointers.

Related Topics

Related Keywords

- top_method_pointer : Identify which method leads the Dakota study
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method

- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name_list :
- model_pointer_list :
- method_pointer_list :
- global_method_name : Specify the global method
- global_method_pointer : Specify the global method
- local_method_name : Specify the local method
- local_method_pointer : Specify the local method
- method_name_list :
- model_pointer_list :
- method_pointer_list :
- model_pointer : Specify which model block will be used by a method

- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- model_pointer : Specify which model block will be used by a method
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- model_pointer : Specify which model block will be used by a method
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method

- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- optional_interface_pointer :
- optional_interface_responses_pointer :
- sub_method_pointer : The `sub_method_pointer` specifies the method block for the sub-iterator
- responses_pointer : Specify which reponses block will be used by this model block
- interface_pointer : Interface block pointer for the single model type
- dace_method_pointer : Specify a method to gather training data
- actual_model_pointer : Pointer to specify a ”truth” model, from which to construct a surrogate
- actual_model_pointer : Pointer to specify a ”truth” model, from which to construct a surrogate
- variables_pointer : Specify which variables block will be included with this model block
- id_variables : Give the block a name, in case of multiple variables blocks

4.2.3 data_import_capabilities

Description

empty

Related Topics

Related Keywords

4.2.4 dakota_output

Description

empty

Related Topics

Related Keywords

- graphics : Display a 2D graphics window of variables and responses
- output_precision : Control the output precision
- results_output : (Experimental) Write a summary file containing the final results
- results_output_file : The base file name of the results file
- tabular_graphics_data : Write a tabular results file with variable and response history
- tabular_graphics_file : The name of the tabular data output file
- output : Control how much method information is written to the screen and output file

4.2.5 file_formats

Description

empty

Related Topics

Related Keywords

- aprepro : Write parameters files in APREPRO syntax
- aprepro : Write parameters files in APREPRO syntax
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format

- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format

- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format
- annotated : (Default) Specified file is in the annotated format
- freeform : Specified file is in the freeform format

4.3 dakota_concepts

Description

Miscallaneous concepts related to Dakota operation

Related Topics

- method_independent_controls
- block
- strategies
- command_line_options
- restarts
- pointers

Related Keywords

4.3.1 method_independent_controls

Description

The `<method independent controls>` are those controls which are valid for a variety of methods. In some cases, these controls are abstractions which may have slightly different implementations from one method to the next. While each of these controls is not valid for every method, the controls are valid for enough methods that it was reasonable to consolidate the specifications.

Related Topics

- linear_constraints

Related Keywords

- constraint_tolerance : The maximum allowable value of constraint violation still considered to be feasible
- convergence_tolerance : Stopping criterion based on convergence of the objective function
- final_solutions : Number of designs returned as the best solutions
- id_method : Give the block a name, in case of multiple method blocks
- max_function_evaluations : Stopping criteria based on number of function evaluations
- max_iterations : Stopping criteria based on number of iterations
- output : Control how much method information is written to the screen and output file
- scaling : Turn on scaling for variables, responses, and constraints
- speculative : Compute speculative gradients

4.3.2 linear_constraints

Description

Many methods use linear equality or inequality constraints.

Related Topics

Related Keywords

- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint

- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint

- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities

- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint

- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint

- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities

- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint

- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint
- linear_equality_constraint_matrix : Define coefficients of the linear equalities
- linear_equality_scale_types : Specify how each linear equality constraint is scaled
- linear_equality_scales : Define the characteristic values to scale linear equalities
- linear_equality_targets : Define target values for the linear equality constraints
- linear_inequality_constraint_matrix : Define coefficients of the linear inequality constraints
- linear_inequality_lower_bounds : Define lower bounds for the linear inequality constraint
- linear_inequality_scale_types : Specify how each linear inequality constraint is scaled
- linear_inequality_scales : Define the characteristic values to scale linear inequalities
- linear_inequality_upper_bounds : Define upper bounds for the linear inequality constraint

4.3.3 block

Description

A block is the highest level of keyword organization in Dakota. There are currently 6 blocks in the Dakota input spec:

Related Topics

- block_identifier
- block_pointer

Related Keywords

- environment : Top-level settings for Dakota execution
- interface : Specifies how function evaluations will be performed in order to map the variables into the responses.
- method : Specifies the name and behavior of a method
- model : Specifies how variables are mapped into a set of responses
- responses : Description of the model output data returned to Dakota upon evaluation of an interface.
- variables : Specifies the parameter set to be iterated by a particular method.

block_identifier

Description

empty

Related Topics

Related Keywords

- id_interface : Give the block a name, in case of multiple interface blocks
- id_method : Give the block a name, in case of multiple method blocks
- id_model : Give the model block an identifying name, in case of multiple model blocks
- id_responses : Name the response block, when there are multiple responses blocks

block_pointer

Description

See block_pointer for details about pointers.

Related Topics

Related Keywords

- top_method_pointer : Identify which method leads the Dakota study
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method

- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name_list :
- model_pointer_list :
- method_pointer_list :
- global_method_name : Specify the global method
- global_method_pointer : Specify the global method
- local_method_name : Specify the local method
- local_method_pointer : Specify the local method
- method_name_list :
- model_pointer_list :
- method_pointer_list :
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- model_pointer : Specify which model block will be used by a method
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method

- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- model_pointer : Specify which model block will be used by a method
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- optional_interface_pointer :
- optional_interface_responses_pointer :
- sub_method_pointer : The `sub_method_pointer` specifies the method block for the sub-iterator
- responses_pointer : Specify which reponses block will be used by this model block
- interface_pointer : Interface block pointer for the single model type

- dace_method_pointer : Specify a method to gather training data
- actual_model_pointer : Pointer to specify a ”truth” model, from which to construct a surrogate
- actual_model_pointer : Pointer to specify a ”truth” model, from which to construct a surrogate
- variables_pointer : Specify which variables block will be included with this model block
- id_variables : Give the block a name, in case of multiple variables blocks

4.3.4 strategies

Description

empty

Related Topics

- advanced_strategies

Related Keywords

4.3.5 advanced_strategies

Description

empty

Related Topics

Related Keywords

4.3.6 command_line_options

Description

empty

Related Topics

Related Keywords

4.3.7 restarts

Description

empty

Related Topics

Related Keywords

4.3.8 pointers

Description

For all pointer specifications, if a pointer string is specified and no corresponding id string is available, Dakota will exit with an error message.

If the pointer is optional and no pointer string is specified, then the last specification parsed will be used.

It is appropriate to omit optional cross-referencing whenever the relationships are unambiguous due to the presence of only one specification.

Related Topics

- block_pointer
- objective_function_pointer

Related Keywords

4.3.9 block_pointer

Description

See block_pointer for details about pointers.

Related Topics

Related Keywords

- top_method_pointer : Identify which method leads the Dakota study
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method

- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name_list :
- model_pointer_list :
- method_pointer_list :
- global_method_name : Specify the global method
- global_method_pointer : Specify the global method
- local_method_name : Specify the local method
- local_method_pointer : Specify the local method
- method_name_list :
- model_pointer_list :
- method_pointer_list :

- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- model_pointer : Specify which model block will be used by a method
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- model_pointer : Specify which model block will be used by a method
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method

- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- method_name : Name of a sub-method to apply to surrogate model
- method_pointer : Pointer to sub-method to apply to surrogate
- model_pointer : Specify which model block will be used by a method
- model_pointer : Specify which model block will be used by a method
- optional_interface_pointer :
- optional_interface_responses_pointer :
- sub_method_pointer : The `sub_method_pointer` specifies the method block for the sub-iterator
- responses_pointer : Specify which reponses block will be used by this model block
- interface_pointer : Interface block pointer for the single model type
- dace_method_pointer : Specify a method to gather training data
- actual_model_pointer : Pointer to specify a ”truth” model, from which to construct a surrogate
- actual_model_pointer : Pointer to specify a ”truth” model, from which to construct a surrogate
- variables_pointer : Specify which variables block will be included with this model block
- id_variables : Give the block a name, in case of multiple variables blocks

4.3.10 objective_function_pointer

Description

See block_pointer for details about pointers.

Related Topics

Related Keywords

- id_analytic_gradients : Identify which analytical gradient corresponds to which response
- id_numerical_gradients : Identify which numerical gradient corresponds to which response
- id_analytic_hessians : Identify which analytical Hessian corresponds to which response
- id_numerical_hessians : Identify which numerical-Hessian corresponds to which response
- id_quasi_hessians : Identify which quasi-Hessian corresponds to which response

4.4 models

Description

Keywords and Concepts relating to the `model` block

Related Topics

- surrogate_models
- recast_models
- multifidelity_models
- reduced_order_models
- nested_models
- advanced_model_recursion

Related Keywords

4.4.1 surrogate_models

Description

empty

Related Topics

- surrogate_based_optimization_methods

Related Keywords

4.4.2 surrogate_based_optimization_methods

Description

empty

Related Topics

Related Keywords

- efficient_global : Global Surrogate Based Optimization, a.k.a. EGO
- surrogate_based_global : Global Surrogate Based Optimization
- surrogate_based_local : Local Surrogate Based Optimization

4.4.3 recast_models

Description

empty

Related Topics

Related Keywords

4.4.4 multifidelity_models

Description

empty

Related Topics

Related Keywords

4.4.5 reduced_order_models

Description

empty

Related Topics

Related Keywords

4.4.6 nested_models

Description

empty

Related Topics

Related Keywords

4.4.7 advanced_model_recursion

Description

empty

Related Topics

- hybrid_and_recursions_logic

Related Keywords

hybrid_and_recursions_logic

Description

empty

Related Topics

Related Keywords

4.5 variables_and_responses

Description

Keywords and concepts relating to the `variables` and `responses` blocks

Related Topics

- variable_domain
- variable_type
- variables_management
- response_types

Related Keywords

4.5.1 variable_domain

Description

Dakota variables can be grouped by their valid domains.

1. Continuous variables are defined by a single interval of real numbers
2. Discrete variables are defined by a finite set of real numbers

Related Topics

- continuous_variables
- discrete_variables

Related Keywords

4.5.2 continuous_variables

Description

empty

Related Topics

Related Keywords

- beta_uncertain : Aleatory uncertain variable - beta
- continuous_design : Continuous design variables, each defined by a single interval

- continuous_interval_uncertain : Epistemic uncertain variable - continuous values within one or more intervals
- continuous_state : Continuous state variables
- discrete_interval_uncertain : Epistemic uncertain variable - one or more intervals
- exponential_uncertain : Aleatory uncertain variable - exponential
- frechet_uncertain : Aleatory uncertain variable - Frechet
- gamma_uncertain : Aleatory uncertain variable - gamma
- gumbel_uncertain : Aleatory uncertain variable - gumbel
- histogram_bin_uncertain : Aleatory uncertain variable - an arbitrary continuous random variable !!!!!!!!!!!!!!!clarify difference between ordinatea and count
- lognormal_uncertain : Aleatory uncertain variable - lognormal
- loguniform_uncertain : Aleatory uncertain variable - loguniform
- normal_uncertain : Aleatory uncertain variable - normal (Gaussian)
- triangular_uncertain : Aleatory uncertain variable - triangular
- uniform_uncertain : Aleatory uncertain variable - uniform
- weibull_uncertain : Aleatory uncertain variable - Weibull

4.5.3 discrete_variables

Description

This page discusses discrete state, design, and uncertain variables.

Variables with the word `discrete` in their name are all specified in similar ways. These include:

1. Set of integers
2. Set of reals
3. Range(s) and each of these is described below.

In addition, some uncertain variables are discrete, real-valued random variables and are specified using parameters. These are described on their individual keyword pages.

Sets

Sets of integers and reals are the same, except that one of them is restricted to integer values.

The variables are specified using three keywords:

- Variable declaration keyword - specifies the number of variables being defined
- `num_set_values` - a list of positive integers specifying how many members each variable contains
  - Length = # of variables
- `set_values` - a list of the permissable integer values in ALL sets, concatenated together.
  - Length = sum of `num_set_values`

  - The order is very important here.
  - The list is partitioned according to the values of `num_set_values`, and each partition is assigned to a variable.

- The ordering of `num_set_values`, and the partitions of `set_values` must match the strings from `descriptions`

Examples are given on the pages:

- discrete_design_set_integer
- discrete_design_set_real
- discrete_uncertain_set_integer
- discrete_uncertain_set_real

Range

For discrete variables defined by range(s), the `lower_bounds` and `upper_bounds` restrict the permisible values. For design variables, this constrains the feasible design space and is frequently used to prevent nonphysical designs. This is a discrete interval variable that may take any integer value within bounds (e.g., [1, 4], allowing values of 1, 2, 3, or 4). For some variable types, each variable is can be defined by multiple ranges.

Examples are given on the pages:

- discrete_interval_uncertain

Related Topics

Related Keywords

- binomial_uncertain : Aleatory uncertain discrete variable - binomial
- discrete_design_range : Discrete design variable - specified using a range of integers
- discrete_design_set_integer : Discrete design variables - specified using lists of integers
- discrete_design_set_real : Discrete design variables - specified using lists of reals
- discrete_state_range : Discrete state variable defined by bounds
- discrete_state_set_integer : Discrete state variables, each defined by a set of permissible integers
- discrete_state_set_real : Discrete state variables, each defined by a set of permissible real numbers
- discrete_uncertain_set_integer : Discrete, epistemic uncertain variable - integers within a set
- discrete_uncertain_set_real : Discrete, epistemic uncertain variable - real numbers within a set
- geometric_uncertain : Aleatory uncertain discrete variable - geometric
- histogram_point_uncertain : Aleatory uncertain discrete variable - arbitrary discrete random variable
- hypergeometric_uncertain : Aleatory uncertain discrete variable - hypergeometric
- negative_binomial_uncertain : Aleatory uncertain discrete variable - negative binomial
- poisson_uncertain : Aleatory uncertain discrete variable - Poisson

4.5.4 variable_type

Description

Dakota variables can be grouped by their type or usage. Each type is intended to be used with a set of methods, i.e.: design variables used with optimization methods. However, the user can also change how methods use the different variable types - see active.

Related Topics

- design_variables
- aleatory_uncertain_variables
- epistemic_uncertain_variables
- state_variables

Related Keywords

4.5.5 design_variables

Description

Design variables are those variables which are modified for the purposes of computing an optimal design.

The most common type of design variables encountered in engineering applications are of the continuous type. These variables may assume any real value within their bounds. All but a handful of the optimization algorithms in Dakota support continuous design variables exclusively.

Related Topics

Related Keywords

- continuous_design : Continuous design variables, each defined by a single interval
- discrete_design_range : Discrete design variable - specified using a range of integers
- discrete_design_set_integer : Discrete design variables - specified using lists of integers
- discrete_design_set_real : Discrete design variables - specified using lists of reals

4.5.6 aleatory_uncertain_variables

Description

Aleatory uncertainty is also known as inherent variability, irreducible uncertainty, or randomness.

Aleatory uncertainty is predominantly charaterized using probability theory. This is the only option implemented in Dakota.

Related Topics

Related Keywords

- beta_uncertain : Aleatory uncertain variable - beta
- binomial_uncertain : Aleatory uncertain discrete variable - binomial
- exponential_uncertain : Aleatory uncertain variable - exponential
- frechet_uncertain : Aleatory uncertain variable - Frechet
- gamma_uncertain : Aleatory uncertain variable - gamma
- geometric_uncertain : Aleatory uncertain discrete variable - geometric
- gumbel_uncertain : Aleatory uncertain variable - gumbel
- histogram_bin_uncertain : Aleatory uncertain variable - an arbitrary continuous random variable !!!!!!!!!!!!!!clarify difference between ordinatea and count
- histogram_point_uncertain : Aleatory uncertain discrete variable - arbitrary discrete random variable
- hypergeometric_uncertain : Aleatory uncertain discrete variable - hypergeometric
- lognormal_uncertain : Aleatory uncertain variable - lognormal
- loguniform_uncertain : Aleatory uncertain variable - loguniform
- negative_binomial_uncertain : Aleatory uncertain discrete variable - negative binomial
- normal_uncertain : Aleatory uncertain variable - normal (Gaussian)
- poisson_uncertain : Aleatory uncertain discrete variable - Poisson
- triangular_uncertain : Aleatory uncertain variable - triangular
- uniform_uncertain : Aleatory uncertain variable - uniform
- weibull_uncertain : Aleatory uncertain variable - Weibull

4.5.7 epistemic_uncertain_variables

Description

Epistemic uncertainty is uncertainty due to lack of knowledge.

In Dakota, epistemic uncertainty is characterized by interval analysis or the Dempster-Shafer theory of evidence.

Note that epistemic uncertainty can also be modeled with probability density functions - similarly to aleatory uncertainty Dakota does not support this capability.

Related Topics

Related Keywords

- continuous_interval_uncertain : Epistemic uncertain variable - continuous values within one or more intervals
- discrete_interval_uncertain : Epistemic uncertain variable - one or more intervals
- discrete_uncertain_set_integer : Discrete, epistemic uncertain variable - integers within a set
- discrete_uncertain_set_real : Discrete, epistemic uncertain variable - real numbers within a set

4.5.8 state_variables

Description

State variables provide a convenient mechanism for managing additional model parameterizations such as mesh density, simulation convergence tolerances, and time step controls.

Only parameter studies and design of experiments methods will iterate on state variables.

The `initial_state` is used as the only value for the state variable for all other methods, unless `active state` is invoked.

If a method iterates on a state variable, the variable is treated as a design variable with the given bounds, or as a uniform uncertain variable with the given bounds.

If the state variable is defined only by its bounds, and the method does NOT iterate on state variables, then the `initial_state` must be inferred.

Related Topics

Related Keywords

- continuous_state : Continuous state variables
- discrete_state_range : Discrete state variable defined by bounds
- discrete_state_set_integer : Discrete state variables, each defined by a set of permissible integers
- discrete_state_set_real : Discrete state variables, each defined by a set of permissible real numbers

4.5.9 variables_management

Description

empty

Related Topics

Related Keywords

4.5.10 response_types

Description

Dakota responses are grouped by type or usage. Each type is intended to be used with a set of methods, i.e.: objective functions used with optimization methods.

Related Topics

Related Keywords

4.6 interface

Description

Keywords and Concepts relating to the `interface` block, which is used to connect Dakota to external analysis codes (simulations, etc.)

Related Topics

- simulation_file_management
- workflow_management
- advanced_simulation_interfaces

Related Keywords

4.6.1 simulation_file_management

Description

empty

Related Topics

Related Keywords

4.6.2 workflow_management

Description

empty

Related Topics

Related Keywords

4.6.3 advanced_simulation_interfaces

Description

empty

Related Topics

- simulation_failure
- concurrency_and_parallelism

Related Keywords

simulation_failure

Description

empty

Related Topics

Related Keywords

concurrency_and_parallelism

Description

empty

Related Topics

Related Keywords

- processors_per_analysis : Specify the number of processors per analysis when Dakota is run in parallel
- analysis_scheduling : Specify the scheduling of concurrent analyses when Dakota is run in parallel
- master : Specify a dedicated master partition for parallel analysis scheduling
- peer : Specify a peer partition for parallel analysis scheduling
- analysis_servers : Specify the number of analysis servers when Dakota is run in parallel
- asynchronous : Specify analysis driver concurrency, when Dakota is run in serial
- analysis_concurrency : Limit the number of analysis drivers within an evaluation that Dakota will schedule
- evaluation_concurrency : Determine how many concurrent evaluations Dakota will schedule
- local_evaluation_scheduling : Control how local asynchronous jobs are scheduled
- master : Specify a dedicated master partition for parallel evaluation scheduling
- peer : Specify a peer partition for parallel evaluation scheduling
- dynamic : Specify dynamic scheduling in a peer partition when Dakota is run in parallel.
- static : Specify static scheduling in a peer partition when Dakota is run in parallel.
- evaluation_servers : Specify the number of evaluation servers when Dakota is run in parallel
- processors_per_evaluation : Specify the number of processors per evaluation server when Dakota is run in parallel
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- master : Specify a dedicated master partition for parallel iterator scheduling
- peer : Specify a peer partition for parallel iterator scheduling

- iterator_servers : Specify the number of iterator servers when Dakota is run in parallel
- processors_per_iterator : Specify the number of processors per iterator server when Dakota is run in parallel
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- master : Specify a dedicated master partition for parallel iterator scheduling
- peer : Specify a peer partition for parallel iterator scheduling
- iterator_servers : Specify the number of iterator servers when Dakota is run in parallel
- processors_per_iterator : Specify the number of processors per iterator server when Dakota is run in parallel
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- master : Specify a dedicated master partition for parallel iterator scheduling
- peer : Specify a peer partition for parallel iterator scheduling
- iterator_servers : Specify the number of iterator servers when Dakota is run in parallel
- processors_per_iterator : Specify the number of processors per iterator server when Dakota is run in parallel
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- master : Specify a dedicated master partition for parallel iterator scheduling
- peer : Specify a peer partition for parallel iterator scheduling
- iterator_servers : Specify the number of iterator servers when Dakota is run in parallel
- processors_per_iterator : Specify the number of processors per iterator server when Dakota is run in parallel

4.7 methods

Description

Keywords and Concepts relating to the `method` block, including discussion of the different methods and algorithms availabe in Dakota

Related Topics

- parameter_studies
- sensitivity_analysis_and_design_of_experiments
- uncertainty_quantification
- optimization_and_calibration

Related Keywords

4.7.1 parameter_studies

Description

Parameter studies employ deterministic designs to explore the effect of parametric changes within simulation models, yielding one form of sensitivity analysis. They can help assess simulation characteristics such as smoothness, multi-modality, robustness, and nonlinearity, which affect the choice of algorithms and controls in follow-on optimization and UQ studies.

Dakota's parameter study methods compute response data sets at a selection of points in the parameter space. These points may be specified as a vector, a list, a set of centered vectors, or a multi-dimensional grid. Capability overviews and examples of the different types of parameter studies are provided in the Users Manual [[4] "Adams et al., 2010"].

With the exception of output verbosity (a setting of `silent` will suppress some parameter study diagnostic output), Dakota's parameter study methods do not make use of the method independent controls. Therefore, the parameter study documentation which follows is limited to the method dependent controls for the vector, list, centered, and multidimensional parameter study methods.

Related Topics

Related Keywords

- centered_parameter_study : Parameter study - points moving out from a center point
- list_parameter_study : Parameter study specified by a list of variable values
- multidim_parameter_study : Parameter study - make a full factorial grid of study points
- vector_parameter_study : Parameter study - samples points along a vector

4.7.2 sensitivity_analysis_and_design_of_experiments

Description

empty

Related Topics

- design_and_analysis_of_computer_experiments
- sampling

Related Keywords

4.7.3 design_and_analysis_of_computer_experiments

Description

Design and Analysis of Computer Experiments (DACE) methods compute response data sets at a selection of points in the parameter space. Three libraries are provided for performing these studies: DDACE, FSUDace, and PSUADE. The design of experiments methods do not currently make use of any of the method independent controls.

Related Topics

Related Keywords

- dace : Design and Analysis of Computer Experiments
- fsu_cvt : Design of Computer Experiments - Centroidal Voronoi Tessellation
- fsu_quasi_mc : Design of Computer Experiments - Quasi-Monte Carlo sampling
- psuade_moat : Morris One-at-a-Time

4.7.4 sampling

Description

Sampling techniques are selected using the `sampling` method selection. This method generates sets of samples according to the probability distributions of the uncertain variables and maps them into corresponding sets of response functions, where the number of samples is specified by the `samples` integer specification. Means, standard deviations, coefficients of variation (COVs), and 95% confidence intervals are computed for the response functions. Probabilities and reliabilities may be computed for `response_levels` specifications, and response levels may be computed for either `probability_levels` or `reliability_levels` specifications (refer to the Method Commands chapter in the Dakota Reference Manual[5] for additional information).

Currently, traditional Monte Carlo (MC) and Latin hypercube sampling (LHS) are supported by Dakota and are chosen by specifying `sample_type` as `random` or `lhs`. In Monte Carlo sampling, the samples are selected randomly according to the user-specified probability distributions. Latin hypercube sampling is a stratified sampling technique for which the range of each uncertain variable is divided into N_s segments of equal probability, where N_s is the number of samples requested. The relative lengths of the segments are determined by the nature of the specified probability distribution (e.g., uniform has segments of equal width, normal has small segments near the mean and larger segments in the tails). For each of the uncertain variables, a sample is selected randomly from each of these equal probability segments. These N_s values for each of the individual parameters are then combined in a shuffling operation to create a set of N_s parameter vectors with a specified correlation structure. A feature of the resulting sample set is that *every row and column in the hypercube of partitions has exactly one sample*. Since the total number of samples is exactly equal to the number of partitions used for each uncertain variable, an arbitrary number of desired samples is easily accommodated (as compared to less flexible approaches in which the total number of samples is a product or exponential function of the number of intervals for each variable, i.e., many classical design of experiments methods).

Advantages of sampling-based methods include their relatively simple implementation and their independence from the scientific disciplines involved in the analysis. The main drawback of these techniques is the large number of function evaluations needed to generate converged statistics, which can render such an analysis computationally very expensive, if not intractable, for real-world engineering applications. LHS techniques, in general, require fewer samples than traditional Monte Carlo for the same accuracy in statistics, but they still can be prohibitively expensive. For further information on the method and its relationship to other sampling techniques, one is referred to the works by McKay, et al.[60], Iman and Shortencarier[53], and Helton and Davis[46]. Note that under certain separability conditions associated with the function to be sampled, Latin hypercube sampling provides a more accurate estimate of the mean value than does random sampling. That is, given an equal number of samples, the LHS estimate of the mean will have less variance than the mean value obtained through random sampling.

Related Topics

Related Keywords

- importance_sampling : Importance sampling
- sampling : Uncertainty quantification based on random sampling of the variables

4.7.5 uncertainty_quantification

Description

Dakota provides a variety of methods for propagating both aleatory and epistemic uncertainty.

At a high level, uncertainty quantification (UQ) or nondeterministic analysis is the process of characterizing input uncertainties, forward propagating these uncertainties through a computational model, and performing statistical or interval assessments on the resulting responses. This process determines the effect of uncertainties and assumptions on model outputs or results. In Dakota, uncertainty quantification methods specifically focus on the forward propagation part of the process, where probabilistic or interval information on parametric inputs are mapped through the computational model to assess statistics or intervals on outputs. For an overview of these approaches for engineering applications, consult[42].

UQ is related to sensitivity analysis in that the common goal is to gain an understanding of how variations in the parameters affect the response functions of the engineering design problem. However, for UQ, some or all of the components of the parameter vector, are considered to be uncertain as specified by particular probability distributions (e.g., normal, exponential, extreme value), or other uncertainty structures. By assigning specific distributional structure to the inputs, distributional structure for the outputs (i.e, response statistics) can be inferred. This migrates from an analysis that is more $\{qualitative\}$ in nature, in the case of sensitivity analysis, to an analysis that is more rigorously $\{quantitative\}$.

UQ methods are often distinguished by their ability to propagate aleatory or epistemic input uncertainty characterizations, where aleatory uncertainties are irreducible variabilities inherent in nature and epistemic uncertainties are reducible uncertainties resulting from a lack of knowledge. Since sufficient data is generally available for aleatory uncertainties, probabilistic methods are commonly used for computing response distribution statistics based on input probability distribution specifications. Conversely, for epistemic uncertainties, any use of probability distributions is based on subjective knowledge rather than objective data, and we may alternatively explore nonprobabilistic methods based on interval specifications.

Dakota contains capabilities for performing nondeterministic analysis with both types of input uncertainty. These UQ methods have been developed by Sandia Labs, in conjunction with collaborators in academia[30],[31],[20],[80].

The aleatory UQ methods in Dakota include various sampling-based approaches (e.g., Monte Carlo and Latin Hypercube sampling), local and global reliability methods, and stochastic expansion (polynomial chaos expansions and stochastic collocation) approaches. The epistemic UQ methods include local and global interval analysis and Dempster-Shafer evidence theory. These are summarized below and then described in more depth in subsequent sections of this chapter. Dakota additionally supports mixed aleatory/epistemic UQ via interval-valued probability, second-order probability, and Dempster-Shafer theory of evidence. These involve advanced model recursions and are described in Section.

Dakota contains capabilities for performing nondeterministic analysis with both types of input uncertainty. These UQ methods have been developed by Sandia Labs, in conjunction with collaborators in academia[30],[31],[20],[80].

The aleatory UQ methods in Dakota include various sampling-based approaches (e.g., Monte Carlo and Latin Hypercube sampling), local and global reliability methods, and stochastic expansion (polynomial chaos expansions and stochastic collocation) approaches. The epistemic UQ methods include local and global interval analysis and Dempster-Shafer evidence theory. These are summarized below and then described in more depth in subsequent sections of this chapter. Dakota additionally supports mixed aleatory/epistemic UQ via interval-valued

probability, second-order probability, and Dempster-Shafer theory of evidence. These involve advanced model recursions and are described in Section.

The choice of uncertainty quantification method depends on how the input uncertainty is characterized, the computational budget, and the desired output accuracy. The recommendations for UQ methods are summarized in Tableand are discussed in the remainder of the section.

TODO: Put table in Doxygen if still needed

Related Topics

- aleatory_uncertainty_quantification_methods
- epistemic_uncertainty_quantification_methods

Related Keywords

- adaptive_sampling : Build a GP surrogate and refine it adaptively
- global_interval_est : Interval analysis using global optimization methods
- global_reliability : Global reliability methods
- gpais : Gaussian Process Adaptive Importance Sampling
- importance_sampling : Importance sampling
- local_interval_est : Interval analysis using local optimization
- local_reliability : Local reliability method
- mpp_search : Specify which MPP search option to use
- pof_darts : Probability-of-Failure (POF) darts is a novel method for estimating the probability of failure based on random sphere-packing.
- sampling : Uncertainty quantification based on random sampling of the variables

4.7.6 aleatory_uncertainty_quantification_methods

Description

Aleatory uncertainty is also known as inherent variability, irreducible uncertainty, or randomness.

Aleatory uncertainty is typically charaterized using probability theory.

Related Topics

- sampling
- reliability_methods
- stochastic_expansion_methods

Related Keywords

- importance_sampling : Importance sampling

sampling

Description

Sampling techniques are selected using the `sampling` method selection. This method generates sets of samples according to the probability distributions of the uncertain variables and maps them into corresponding sets of response functions, where the number of samples is specified by the `samples` integer specification. Means, standard deviations, coefficients of variation (COVs), and 95% confidence intervals are computed for the response functions. Probabilities and reliabilities may be computed for `response_levels` specifications, and response levels may be computed for either `probability_levels` or `reliability_levels` specifications (refer to the Method Commands chapter in the Dakota Reference Manual[5] for additional information).

Currently, traditional Monte Carlo (MC) and Latin hypercube sampling (LHS) are supported by Dakota and are chosen by specifying `sample_type` as `random` or `lhs`. In Monte Carlo sampling, the samples are selected randomly according to the user-specified probability distributions. Latin hypercube sampling is a stratified sampling technique for which the range of each uncertain variable is divided into N_s segments of equal probability, where N_s is the number of samples requested. The relative lengths of the segments are determined by the nature of the specified probability distribution (e.g., uniform has segments of equal width, normal has small segments near the mean and larger segments in the tails). For each of the uncertain variables, a sample is selected randomly from each of these equal probability segments. These N_s values for each of the individual parameters are then combined in a shuffling operation to create a set of N_s parameter vectors with a specified correlation structure. A feature of the resulting sample set is that *every row and column in the hypercube of partitions has exactly one sample*. Since the total number of samples is exactly equal to the number of partitions used for each uncertain variable, an arbitrary number of desired samples is easily accommodated (as compared to less flexible approaches in which the total number of samples is a product or exponential function of the number of intervals for each variable, i.e., many classical design of experiments methods).

Advantages of sampling-based methods include their relatively simple implementation and their independence from the scientific disciplines involved in the analysis. The main drawback of these techniques is the large number of function evaluations needed to generate converged statistics, which can render such an analysis computationally very expensive, if not intractable, for real-world engineering applications. LHS techniques, in general, require fewer samples than traditional Monte Carlo for the same accuracy in statistics, but they still can be prohibitively expensive. For further information on the method and its relationship to other sampling techniques, one is referred to the works by McKay, et al.[60], Iman and Shortencarier[53], and Helton and Davis[46]. Note that under certain separability conditions associated with the function to be sampled, Latin hypercube sampling provides a more accurate estimate of the mean value than does random sampling. That is, given an equal number of samples, the LHS estimate of the mean will have less variance than the mean value obtained through random sampling.

Related Topics

Related Keywords

- importance_sampling : Importance sampling
- sampling : Uncertainty quantification based on random sampling of the variables

reliability_methods

Description

Reliability methods provide an alternative approach to uncertainty quantification which can be less computationally demanding than sampling techniques. Reliability methods for uncertainty quantification are based on

probabilistic approaches that compute approximate response function distribution statistics based on specified uncertain variable distributions. These response statistics include response mean, response standard deviation, and cumulative or complementary cumulative distribution functions (CDF/CCDF). These methods are often more efficient at computing statistics in the tails of the response distributions (events with low probability) than sampling based approaches since the number of samples required to resolve a low probability can be prohibitive.

The methods all answer the fundamental question: "Given a set of uncertain input variables, $\mathbf{X}$, and a scalar response function, g, what is the probability that the response function is below or above a certain level, $\bar{z}$?" The former can be written as $P[g(\mathbf{X}) \leq \bar{z}] = F_g(\bar{z})$ where $F_g(\bar{z})$ is the cumulative distribution function (CDF) of the uncertain response $g(\mathbf{X})$ over a set of response levels. The latter can be written as $P[g(\mathbf{X}) > \bar{z}]$ and defines the complementary cumulative distribution function (CCDF).

This probability calculation involves a multi-dimensional integral over an irregularly shaped domain of interest, $\mathbf{D}$, where $g(\mathbf{X}) < z$ as displayed in Figure figUQ05 for the case of two variables. The reliability methods all involve the transformation of the user-specified uncertain variables, $\mathbf{X}$, with probability density function, $p(x_1, x_2)$, which can be non-normal and correlated, to a space of independent Gaussian random variables, $\mathbf{u}$, possessing a mean value of zero and unit variance (i.e., standard normal variables). The region of interest, $\mathbf{D}$, is also mapped to the transformed space to yield, $\mathbf{D_u}$, where $g(\mathbf{U}) < z$ as shown in Figure figUQ06. The Nataf transformation[16], which is identical to the Rosenblatt transformation[73] in the case of independent random variables, is used in Dakota to accomplish this mapping. This transformation is performed to make the probability calculation more tractable. In the transformed space, probability contours are circular in nature as shown in Figure figUQ06 unlike in the original uncertain variable space, Figure figUQ05 . Also, the multi-dimensional integrals can be approximated by simple functions of a single parameter, β, called the reliability index. β is the minimum Euclidean distance from the origin in the transformed space to the response surface. This point is also known as the most probable point (MPP) of failure. Note, however, the methodology is equally applicable for generic functions, not simply those corresponding to failure criteria; this nomenclature is due to the origin of these methods within the disciplines of structural safety and reliability. Note that there are local and global reliability methods. The majority of the methods available are local, meaning that a local optimization formulation is used to locate one MPP. In contrast, global methods can find multiple MPPs if they exist.

Related Topics

Related Keywords

- global_reliability : Global reliability methods
- u_gaussian_process : Create GP surrogate in u-space
- x_gaussian_process : Create GP surrogate in x-space
- local_reliability : Local reliability method
- mpp_search : Specify which MPP search option to use
- integration :
- first_order :
- second_order :
- no_approx :
- u_taylor_mean :
- u_taylor_mpp :

- u_two_point :
- x_taylor_mean :
- x_taylor_mpp :
- x_two_point :

stochastic_expansion_methods

Description

The development of these techniques mirrors that of deterministic finite element analysis utilizing the notions of projection, orthogonality, and weak convergence[30],[31]. Rather than estimating point probabilities, they form an approximation to the functional relationship between response functions and their random inputs, which provides a more complete uncertainty representation for use in multi-code simulations. Expansion methods include polynomial chaos expansions (PCE), which employ multivariate orthogonal polynomials that are tailored to representing particular input probability distributions, and stochastic collocation (SC), which employs multivariate interpolation polynomials. For PCE, expansion coefficients may be evaluated using a spectral projection approach (based on sampling, tensor-product quadrature, Smolyak sparse grid, or cubature methods for numerical integration) or a regression approach (least squares or compressive sensing). For SC, interpolants are formed over tensor-product or sparse grids and may be local or global, value-based or gradient-enhanced, and nodal or hierarchical. In global value-based cases (Lagrange polynomials), the barycentric formulation is used[9],[57],[49] to improve numerical efficiency and stability. Both sets of methods provide analytic response moments and variance-based metrics; however, CDF/CCDF probabilities are evaluated numerically by sampling on the expansion.

Related Topics

Related Keywords

4.7.7 epistemic_uncertainty_quantification_methods

Description

Epistemic uncertainty is uncertainty due to lack of knowledge.

In Dakota, epistemic uncertainty analysis is performed using interval analysis or Dempster-Shafer theory of evidence.

Note that epistemic uncertainty can also be modeled probabilistically. It would be more accurate to call this class of method, non-probabilistic uncertainty quantification, but the name persists for historical reasons.

Related Topics

- interval_estimation
- evidence_theory

Related Keywords

- global_evidence : Evidence theory with evidence measures computed with global optimization methods
- global_interval_est : Interval analysis using global optimization methods
- local_evidence : Evidence theory with evidence measures computed with local optimization methods

- local_interval_est : Interval analysis using local optimization

interval_estimation

Description

In interval analysis, one assumes that nothing is known about an epistemic uncertain variable except that its value lies somewhere within an interval. In this situation, it is NOT assumed that the value has a uniform probability of occuring within the interval. Instead, the interpretation is that any value within the interval is a possible value or a potential realization of that variable. In interval analysis, the uncertainty quantification problem is one of determining the resulting bounds on the output (defining the output interval) given interval bounds on the inputs. Again, any output response that falls within the output interval is a possible output with no frequency information assigned to it.

We have the capability to perform interval analysis using either `global_interval_est` or `local_interval_est`. In the global approach, one uses either a global optimization method or a sampling method to assess the bounds. `global_interval_est` allows the user to specify either `lhs`, which performs Latin Hypercube Sampling and takes the minimum and maximum of the samples as the bounds (no optimization is performed) or `ego`. In the case of `ego`, the efficient global optimization method is used to calculate bounds. The ego method is described in Section . If the problem is amenable to local optimization methods (e.g. can provide derivatives or use finite difference method to calculate derivatives), then one can use local methods to calculate these bounds. `local_interval_est` allows the user to specify either `sqp` which is sequential quadratic programming, or `nip` which is a nonlinear interior point method.

Note that when performing interval analysis, it is necessary to define interval uncertain variables as described in Section . For interval analysis, one must define only one interval per input variable, in contrast with Dempster-Shafer evidence theory, where an input can have several possible intervals. Interval analysis can be considered a special case of Dempster-Shafer evidence theory where each input is defined by one input interval with a basic probability assignment of one. In Dakota, however, the methods are separate and semantic differences exist in the output presentation. If you are performing a pure interval analysis, we recommend using either `global_interval_est` or `local_interval_est` instead of `global_evidence` or `local_evidence`, for reasons of simplicity. An example of interval estimation is found in the `Dakota/examples/users/cantilever_uq_global_interval.in`, and also in Section .

Note that we have kept separate implementations of interval analysis and Dempster-Shafer evidence theory because our users often want to couple interval analysis on an `outer loop'' with an aleatory, probabilistic analysis on an`inner loop" for nested, second-order probability calculations. See Section for additional details on these nested approaches. These interval methods can also be used as the outer loop within an interval-valued probability analysis for propagating mixed aleatory and epistemic uncertainty – refer to Section for additional details.

Interval analysis is often used to model epistemic uncertainty. In interval analysis, the uncertainty quantification problem is one of determining the resulting bounds on the output (defining the output interval) given interval bounds on the inputs.

We can do interval analysis using either `%global_interval_est` or `local_interval_est`. In the global approach, one uses either a global optimization method or a sampling method to assess the bounds, whereas the local method uses gradient information in a derivative-based optimization approach.

An example of interval estimation is shown in Figure , with example results in Figure . This example is a demonstration of calculating interval bounds for three outputs of the cantilever beam problem. The cantilever beam problem is described in detail in Section . Given input intervals of [1,10] on beam width and beam thickness, we can see that the interval estimate of beam weight is approximately [1,100].

:examples:interval_out

```
-----------------------------------------------------------------
Min and Max estimated values for each response function:
```

```
weight:  Min = 1.0000169352e+00  Max = 9.9999491948e+01
stress:  Min = -9.7749994284e-01  Max = 2.1499428450e+01
displ:  Min = -9.9315672724e-01  Max = 6.7429714485e+01
-----------------------------------------------------------------
```

Related Topics

Related Keywords

- global_interval_est : Interval analysis using global optimization methods
- local_interval_est : Interval analysis using local optimization

evidence_theory

Description

This section discusses Dempster-Shafer evidence theory. In this approach, one does not assign a probability distribution to each uncertain input variable. Rather, one divides each uncertain input variable into one or more intervals. The input parameters are only known to occur within intervals: nothing more is assumed.

Each interval is defined by its upper and lower bounds, and a Basic Probability Assignment (BPA) associated with that interval. The BPA represents a probability of that uncertain variable being located within that interval.

The intervals and BPAs are used to construct uncertainty measures on the outputs called "belief" and "plausibility." Belief represents the smallest possible probability that is consistent with the evidence, while plausibility represents the largest possible probability that is consistent with the evidence. For more information about the Dempster-Shafer theory of evidence, see [67] "Oberkampf and Helton, 2003" and [47] "Helton and Oberkampf, 2004".

Similar to the interval approaches, one may use global or local methods to determine plausbility and belief measures for the outputs.

Usage Notes

Note that to calculate the plausibility and belief cumulative distribution functions, one has to look at all combinations of intervals for the uncertain variables. Within each interval cell combination, the minimum and maximum value of the objective function determine the belief and plausibility, respectively. In terms of implementation, global methods use LHS sampling or global optimization to calculate the minimum and maximum values of the objective function within each interval cell, while local methods use gradient-based optimization methods to calculate these minima and maxima.

Finally, note that many non-deterministic keywords apply to the evidence methods, but one needs to be careful about the interpretation and translate probabilistic measures to epistemic ones. For example, if the user specifies distribution of type complementary, a complementary plausibility and belief function will be generated for the evidence methods (as opposed to a complementary distribution function in the `sampling` case). If the user specifies a set of responses levels, both the belief and plausibility will be calculated for each response level. Likewise, if the user specifies a probability level, the probability level will be interpreted both as a belief and plausibility, and response levels corresponding to the belief and plausibility levels will be calculated. Finally, if generalized reliability levels are specified, either as inputs (`gen_reliability_levels`) or outputs (`response_levels` with `compute gen_reliabilities`), then these are directly converted to/from probability levels and the same probability-based mappings described above are performed.

Related Topics

Related Keywords

- global_evidence : Evidence theory with evidence measures computed with global optimization methods

- local_evidence : Evidence theory with evidence measures computed with local optimization methods

4.7.8 optimization_and_calibration

Description

Optimization algorithms work to minimize (or maximize) an objective function, typically calculated by the user simulation code, subject to constraints on design variables and responses. Available approaches in Dakota include well-tested, proven gradient-based, derivative-free local, and global methods for use in science and engineering design applications. Dakota also offers more advanced algorithms, e.g., to manage multi-objective optimization or perform surrogate-based minimization. This chapter summarizes optimization problem formulation, standard algorithms available in Dakota (mostly through included third-party libraries, see), some advanced capabilities, and offers usage guidelines.

Optimization Formulations

This section provides a basic introduction to the mathematical formulation of optimization, problems. The primary goal of this section is to introduce terms relating to these topics, and is not intended to be a description of theory or numerical algorithms. For further details, consult[7] ,[34],[41],[66], and [85].

A general optimization problem is formulated as follows:

$$\begin{aligned}
\text{minimize:} \quad & f(\mathbf{x}) \\
& \mathbf{x} \in \Re^n \\
\text{subject to:} \quad & \mathbf{g}_L \leq \mathbf{g}(\mathbf{x}) \leq \mathbf{g}_U \\
& \mathbf{h}(\mathbf{x}) = \mathbf{h}_t \\
& \mathbf{a}_L \leq \mathbf{A}_i\mathbf{x} \leq \mathbf{a}_U \\
& \mathbf{A}_e\mathbf{x} = \mathbf{a}_t \\
& \mathbf{x}_L \leq \mathbf{x} \leq \mathbf{x}_U
\end{aligned} \tag{4.1}$$

where vector and matrix terms are marked in bold typeface. In this formulation, $\mathbf{x} = [x_1, x_2, \ldots, x_n]$ is an n-dimensional vector of real-valued *design variables* or *design parameters*. The n-dimensional vectors, $\mathbf{x}_L$ and $\mathbf{x}_U$, are the lower and upper bounds, respectively, on the design parameters. These bounds define the allowable values for the elements of $\mathbf{x}$, and the set of all allowable values is termed the *design space* or the *parameter space*. A *design point* or a *sample point* is a particular set of values within the parameter space.

The optimization goal is to minimize the *objective function*, $f(\mathbf{x})$, while satisfying the constraints. Constraints can be categorized as either linear or nonlinear and as either inequality or equality. The *nonlinear inequality constraints*}, $\mathbf{g}(\mathbf{x})$, are "2-sided," in that they have both lower and upper bounds, $\mathbf{g}_L$ and $\mathbf{g}_U$, respectively. The *nonlinear equality constraints*, $\mathbf{h}(\mathbf{x})$, have target values specified by $\mathbf{h}_t$. The linear inequality constraints create a linear system $\mathbf{A}_i\mathbf{x}$, where $\mathbf{A}_i$ is the coefficient matrix for the linear system. These constraints are also 2-sided as they have lower and upper bounds, $\mathbf{a}_L$ and $\mathbf{a}_U$, respectively. The linear equality constraints create a linear system $\mathbf{A}_e\mathbf{x}$, where $\mathbf{A}_e$ is the coefficient matrix for the linear system and $\mathbf{a}_t$ are the target values. The constraints partition the parameter space into feasible and infeasible regions. A design point is said to be *feasible* if and only if it satisfies all of the constraints. Correspondingly, a design point is said to be *infeasible* if it violates one or more of the constraints.

Many different methods exist to solve the optimization problem given by Equation, all of which iterate on $\mathbf{x}$ in some manner. That is, an initial value for each parameter in $\mathbf{x}$ is chosen, the *response quantities*, $f(\mathbf{x})$, $\mathbf{g}(\mathbf{x})$, $\mathbf{h}(\mathbf{x})$, are computed, often by running a simulation, and some algorithm is applied to generate a new $\mathbf{x}$ that will either reduce the objective function, reduce the amount of infeasibility, or both. To facilitate a general presentation of

these methods, three criteria will be used in the following discussion to differentiate them: optimization problem type, search goal, and search method.

The **optimization problem type** can be characterized both by the types of constraints present in the problem and by the linearity or nonlinearity of the objective and constraint functions. For constraint categorization, a hierarchy of complexity exists for optimization algorithms, ranging from simple bound constraints, through linear constraints, to full nonlinear constraints. By the nature of this increasing complexity, optimization problem categorizations are inclusive of all constraint types up to a particular level of complexity. That is, an *unconstrained problem* has no constraints, a *bound-constrained problem* has only lower and upper bounds on the design parameters, a *linearly-constrained problem* has both linear and bound constraints, and a *nonlinearly-constrained problem* may contain the full range of nonlinear, linear, and bound constraints. If all of the linear and nonlinear constraints are equality constraints, then this is referred to as an *equality-constrained problem*, and if all of the linear and nonlinear constraints are inequality constraints, then this is referred to as an *inequality-constrained problem*. Further categorizations can be made based on the linearity of the objective and constraint functions. A problem where the objective function and all constraints are linear is called a *linear programming (LP) problem*. These types of problems commonly arise in scheduling, logistics, and resource allocation applications. Likewise, a problem where at least some of the objective and constraint functions are nonlinear is called a *nonlinear programming (NLP) problem*. These NLP problems predominate in engineering applications and are the primary focus of Dakota.

The **search goal** refers to the ultimate objective of the optimization algorithm, i.e., either global or local optimization. In *global optimization*, the goal is to find the design point that gives the lowest feasible objective function value over the entire parameter space. In contrast, in *local optimization*, the goal is to find a design point that is lowest relative to a "nearby" region of the parameter space. In almost all cases, global optimization will be more computationally expensive than local optimization. Thus, the user must choose an optimization algorithm with an appropriate search scope that best fits the problem goals and the computational budget.

The **search method** refers to the approach taken in the optimization algorithm to locate a new design point that has a lower objective function or is more feasible than the current design point. The search method can be classified as either *gradient-based* or *nongradient-based*. In a gradient-based algorithm, gradients of the response functions are computed to find the direction of improvement. Gradient-based optimization is the search method that underlies many efficient local optimization methods. However, a drawback to this approach is that gradients can be computationally expensive, inaccurate, or even nonexistent. In such situations, nongradient-based search methods may be useful. There are numerous approaches to nongradient-based optimization. Some of the more well known of these include pattern search methods (nongradient-based local techniques) and genetic algorithms (nongradient-based global techniques).

Because of the computational cost of running simulation models, surrogate-based optimization (SBO) methods are often used to reduce the number of actual simulation runs. In SBO, a surrogate or approximate model is constructed based on a limited number of simulation runs. The optimization is then performed on the surrogate model. Dakota has an extensive framework for managing a variety of local, multipoint, global, and hierarchical surrogates for use in optimization. Finally, sometimes there are multiple objectives that one may want to optimize simultaneously instead of a single scalar objective. In this case, one may employ multi-objective methods that are described in Section.

This overview of optimization approaches underscores that no single optimization method or algorithm works best for all types of optimization problems. Sectionoffers guidelines for choosing a Dakota optimization algorithm best matched to your specific optimization problem.

Constraint Considerations Dakota's input commands permit the user to specify two-sided nonlinear inequality constraints of the form $g_{L_i} \leq g_i(\mathbf{x}) \leq g_{U_i}$, as well as nonlinear equality constraints of the form $h_j(\mathbf{x}) = h_{t_j}$. Some optimizers (e.g., `npsol_`, `optpp_`, `soga`, and `moga` methods) can handle these constraint forms directly, whereas other optimizers (e.g., `asynch_pattern_search`, `dot_`, and `conmin_`, `mesh_adaptive_search`) require Dakota to perform an internal conversion of all constraints to one-sided inequality constraints

of the form $g_i(\mathbf{x}) \leq 0$. In the latter case, the two-sided inequality constraints are treated as $g_i(\mathbf{x}) - g_{U_i} \leq 0$ and $g_{L_i} - g_i(\mathbf{x}) \leq 0$ and the equality constraints are treated as $h_j(\mathbf{x}) - h_{t_j} \leq 0$ and $h_{t_j} - h_j(\mathbf{x}) \leq 0$. The situation is similar for linear constraints: `asynch_pattern_search`, `npsol_`, `optpp_`, `soga`, and `moga` methods support them directly, whereas `dot_` and `conmin_` methods do not. For linear inequalities of the form $a_{L_i} \leq \mathbf{a}_i^T\mathbf{x} \leq a_{U_i}$ and linear equalities of the form $\mathbf{a}_i^T\mathbf{x} = a_{t_j}$, the nonlinear constraint arrays in `dot_` and `conmin_` methods are further augmented to include $\mathbf{a}_i^T\mathbf{x} - a_{U_i} \leq 0$ and $a_{L_i} - \mathbf{a}_i^T\mathbf{x} \leq 0$ in the inequality case and $\mathbf{a}_i^T\mathbf{x} - a_{t_j} \leq 0$ and $a_{t_j} - \mathbf{a}_i^T\mathbf{x} \leq 0$ in the equality case. Awareness of these constraint augmentation procedures can be important for understanding the diagnostic data returned from the `dot_` and `conmin_` methods. Other optimizers fall somewhere in between. `nlpql_` methods support nonlinear equality constraints $h_j(\mathbf{x}) = 0$ and nonlinear one-sided inequalities $g_i(\mathbf{x}) \geq 0$, but does not natively support linear constraints. Constraint mappings are used with NLPQL for both linear and nonlinear cases. Most `coliny_` methods now support two-sided nonlinear inequality constraints and nonlinear constraints with targets, but do not natively support linear constraints.

When gradient and Hessian information is used in the optimization, derivative components are most commonly computed with respect to the active continuous variables, which in this case are the *continuous design variables*. This differs from parameter study methods (for which all continuous variables are active) and from nondeterministic analysis methods (for which the uncertain variables are active). Refer to Sectionfor additional information on derivative components and active continuous variables.

Optimizing with Dakota: Choosing a Method

This section summarizes the optimization methods available in Dakota. We group them according to search method and search goal and establish their relevance to types of problems. For a summary of this discussion, see Section.

Gradient-Based Local Methods Gradient-based optimizers are best suited for efficient navigation to a local minimum in the vicinity of the initial point. They are not intended to find global optima in nonconvex design spaces. For global optimization methods, see. Gradient-based optimization methods are highly efficient, with the best convergence rates of all of the local optimization methods, and are the methods of choice when the problem is smooth, unimodal, and well-behaved. However, these methods can be among the least robust when a problem exhibits nonsmooth, discontinuous, or multimodal behavior. The derivative-free methods described inare more appropriate for problems with these characteristics.

Gradient accuracy is a critical factor for gradient-based optimizers, as inaccurate derivatives will often lead to failures in the search or pre-mature termination of the method. Analytic gradients and Hessians are ideal but often unavailable. If analytic gradient and Hessian information can be provided by an application code, a full Newton method will achieve quadratic convergence rates near the solution. If only gradient information is available and the Hessian information is approximated from an accumulation of gradient data, the superlinear convergence rates can be obtained. It is most often the case for engineering applications, however, that a finite difference method will be used by the optimization algorithm to estimate gradient values. Dakota allows the user to select the step size for these calculations, as well as choose between forward-difference and central-difference algorithms. The finite difference step size should be selected as small as possible, to allow for local accuracy and convergence, but not so small that the steps are "in the noise." This requires an assessment of the local smoothness of the response functions using, for example, a parameter study method. Central differencing will generally produce more reliable gradients than forward differencing but at roughly twice the expense.

Gradient-based methods for nonlinear optimization problems can be described as iterative processes in which a sequence of subproblems, usually which involve an approximation to the full nonlinear problem, are solved until the solution converges to a local optimum of the full problem. The optimization methods available in Dakota fall into several categories, each of which is characterized by the nature of the subproblems solved at each iteration.

Related Topics

- local_optimization_methods
- global_optimization_methods
- bayesian_calibration
- nonlinear_least_squares
- advanced_optimization

Related Keywords

4.7.9 local_optimization_methods

Description

empty

Related Topics

- unconstrained
- constrained
- sequential_quadratic_programming

Related Keywords

- coliny_cobyla : Constrained Optimization BY Linear Approximations (COBYLA)
- nlpql_sqp : Sequential Quadratic Program
- npsol_sqp : Sequential Quadratic Program
- optpp_cg : A conjugate gradient optimization method
- optpp_fd_newton : Finite Difference Newton optimization method
- optpp_g_newton : Newton method based least-squares calbration
- optpp_newton : Newton method based optimization
- optpp_q_newton : Quasi-Newton optimization method

unconstrained

Description

empty

Related Topics

Related Keywords

constrained

Description

empty

Related Topics

Related Keywords

- coliny_cobyla : Constrained Optimization BY Linear Approximations (COBYLA)

sequential_quadratic_programming

Description

Sequential Quadratic Programming (SQP) algorithms are a class of mathematical programming problems used to solve nonlinear optimization problems with nonlinera constraints. These methods are a generalization of Newton's method: each iteration involves minimizing a quadratic model of the problem. These subproblems are formulated as minimizing a quadratic approximation of the Lagrangian subject to linearized constraints. Only gradient information is required; Hessians are approximated by low-rank updates defined by the step taken at each iteration. It is important to note that while the solution found by an SQP method will respect the constraints, the intermediate iterates may not. SQP methods available in Dakota are `dot_sqp`, `nlpql_sqp`, `nlssol_sqp`, and `npsol_sqp`. The particular implementation in `nlpql_sqp` uses a variant with distributed and non-monotone line search. Thus, this variant is designed to be more robust in the presence of inaccurate or noisy gradients common in many engineering applications.

Related Topics

Related Keywords

- nlpql_sqp : Sequential Quadratic Program
- nlssol_sqp : Sequential Quadratic Program for nonlinear least squares
- npsol_sqp : Sequential Quadratic Program

4.7.10 global_optimization_methods

Description

empty

Related Topics

Related Keywords

- asynch_pattern_search : Pattern search, derivative free optimization method
- coliny_direct : DIviding RECTangles method

- coliny_ea : Evolutionary Algorithm
- coliny_pattern_search : Pattern search, derivative free optimization method
- efficient_global : Global Surrogate Based Optimization, a.k.a. EGO
- ncsu_direct : DIviding RECTangles method
- soga : Single-objective Genetic Algorithm (a.k.a Evolutionary Algorithm)

4.7.11 bayesian_calibration

Description

We have three preliminary implementations of Bayesian calibration methods in Dakota, where a "prior distribution" on a parameter is updated through a Bayesian framework involving experimental data and a likelihood function. The theory behind Bayesian methods is best described in other sources [56] and only a brief summary is given here. In Bayesian methods, uncertain parameters are characterized by probability density functions. These probability densities functions define the permissible parameter values - the support, as well as the relative plausibility of each permissible parameter value. In the context of calibration or any inference step, the probability density function that describes knowledge before the incorporation of data is called the prior, $f_{\Theta}(\theta)$.

When data is available, the likelihood function describes how well each parameter value is supported by the data. Bayes Theorem[54], shown in Equation eqnBayesThm, is used for inference: to derive the plausible parameter values, based on the prior probability density and the data d. The result is the posterior parameter density of the parameters $f_{\Theta|D}(\theta|d)$. It is interpreted the same way as the prior, but includes the information derived from the data.

$$f_{\Theta|D}(\theta|d) = \frac{f_{\Theta}(\theta)\,\mathcal{L}(\theta;d)}{f_D(d)}$$

The likelihood function is used to describe how well a model's predictions are supported by the data. The likelihood function can be written generally as:

$$\mathcal{L}(\theta;d) = f(\mathcal{M}(\theta) - d)$$

where θ are the parameters of model $\mathcal{M}$. The function f can greatly influence the results. The specific likelihood functions used in this example were based on Gaussian probability density functions. This means that we assume the difference between the model (e.g. computer simulation) and the experimental observations are Gaussian:

$$d_i = \mathcal{M}(\theta) + \epsilon_i,$$

where ϵ_i is a random variable that can encompass both measurement errors on d_i and modeling errors associated with the simulation $\mathcal{M}(\theta)$. We further assume that all experiments and observations are independent. If we have n observations, the probabilistic model defined by Eq. (TODO BAD REF ref eq:model) results in a likelihood function for θ that is the product of n normal probability density functions as shown in Equation eqnLikelihood.

$$\mathcal{L}(\theta;d) = \prod_{i=1}^{n} \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(d_i - \mathcal{M}(\theta))^2}{2\sigma^2}\right] \mathcal{L}(\theta;d) = \prod_{i=1}^{n} \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(d_i - \mathcal{M}(\theta))^2}{2\sigma^2}\right]$$

Markov Chain Monte Carlo (MCMC) is the standard method used to compute posterior parameter densities, given the observational data and the priors. There are many references that describe the basic algorithm[32], and in addition, the algorithms are an active research area. One variation used in Dakota is DRAM: Delayed Rejection and Adaptive Metropolis[39]. Note that MCMC algorithms take tens or hundreds of thousands of

steps to converge. Since each iteration involves an evaluation of the model $\mathcal{M}(\theta)$, often surrogate models of the simulation model are employed.

As mentioned above, we have three implementations of a Bayesian calibration: one called QUESO, one called DREAM, and one called GPMSA. They are specified with the `bayes_calibration queso` or `bayes_calibration dream` or `bayes_calibration gpmsa`, respectively. The QUESO method uses components from the QUESO library (Quantification of Uncertainty for Estimation, Simulation, and Optimization) developed at The University of Texas at Austin. DREAM uses the DREAM code developed by John Burkardt. It is based on the DiffeRential Evolution Adaptive Metropolis approach which runs multiple different chains simultaneously for global exploration, and automatically tunes the proposal covariance during the process by a self-adaptive randomized subspace sampling [87]. The GPMSA calibration capability uses the GPMSA code developed at Los Alamos National Laboratory.

In the QUESO method, the user can run the MCMC sampling with the simulation model $\mathcal{M}(\theta)$ directly. However, if the model is expensive, we recommend that the user employs a surrogate model (an emulator) because the Monte Carlo Markov Chain will be much faster: the MCMC can generate thousands of samples on the emulator more quickly. One can specify a Gaussian process, a polynomial chaos expansion or a stochastic collocation as the emulator for the `queso` method. The specification details for these are listed in the Reference Manual. One can also specify various settings for the MCMC DRAM sampling: the sampling can use a standard Metropolis--Hastings algorithm or the adaptive Metropolis in which the covariance of the proposal density is updated adaptively. There is also a setting to control the delayed rejection. Finally, there are two scale factors which control the scaling of the problem. The `likelihood_scale` is a number which scales the likelihood by dividing the log of the likelihood (e.g. dividing the sum of squared differences between the experimental data and simulation data or SSE). This is useful for situations with very small likelihoods (e.g. the model is either very far away from the data or there is a lot of data so the likelihood function involves multiplying many likelihoods together, where the SSE term is large and the likelihood becomes very small). In some respects, the `likelihood_scale` can be seen as a normalizing factor for the SSE. If the SSE is large, the likelihood scale should be large. The second factor is a `proposal_covariance_scale` which is a vector that controls the scaling of the proposal covariance in the different input directions. This may be useful when the input variables being calibrated are of different magnitudes: one may want to take a larger step in a direction with a larger magnitude, for example.

For the DREAM method, one can define the number of chains used with `chains`. The total number of generations per chain in DREAM is the number of samples divided by the number of chains. The minimum number of chains is three. The number of chains randomly selected to be used in the crossover each time a crossover occurs is `crossover_chain_pairs`. There is an extra adaptation during burn-in, in which DREAM estimates a distribution of crossover probabilities that favors large jumps over smaller ones in each of the chains. Normalization is required to ensure that all of the input dimensions contribute equally. In this process, a discrete number of candidate points for each crossover value is generated. This parameter is `num_cr`. The `gr_threshold` is the convergence tolerance for the Gelman-Rubin statistic which will govern the convergence of the multiple chain process. The integer `jump_step` forces a long jump every `jump_step` generations. For more details about these parameters, see[87].

GPMSA is another code that provides the capability for Bayesian calibration. A key part of GPMSA is the construction of an emulator from simulation runs collected at various settings of input parameters. The emulator is a statistical model of the system response, and it is used to incorporate the observational data to improve system predictions and constrain or calibrate the unknown parameters. The GPMSA code draws heavily on the theory developed in the seminal Bayesian calibration paper by Kennedy and O'Hagan[56]. The particular approach developed by the Los Alamos group is provided in[48]. GPMSA uses Gaussian process models in the emulation, but the emulator is actually a set of basis functions (e.g. from a singular value decomposition) which have GPs as the coefficients. One major difference between GPMSA and the QUESO implementation in Dakota is that the QUESO implementation does not have an explicit "discrepancy" function δ which models the difference between the simulation and the observational data results in addition to the error term ϵ, but GPMSA has a sophisticated model for the discrepancy term. At this point, the GPMSA implementation in Dakota is an early prototype.

At this point, the GPMSA library is a standalone C++ library which has its own methods to create Gaussian process models, perform MCMC updating, etc. The GPMSA C++ library is an alpha-version and undergoing development, so a user is cautioned to obtain the latest version (e.g. Version-of-the-Day) to have the latest updates. We expect that future Dakota-GPMSA integration will involve more interoperability and sharing of optimization algorithms and surrogate models between Dakota and GPMSA, for example.

We briefly describe the process of running QUESO from Dakota. The user will create a Dakota input file such as the one shown in figUQ18. Note that the method is `bayes_calibration queso`, specifying the QUESO algorithm. The number of samples indicates the number of samples that the MCMC algorithm will take, in this case 5000 (this usually will need to be larger). For this example, we are using the `text_book` analytic example, so we do not need to specify an emulator, but the lines commented out give an idea of the options if the user wanted to specify an emulator. This example is using the full DRAM (delayed rejection adaptive metropolis). The likelihood is scaled, but the proposal covariance is not unless the user uncomments that line. The calibration terms in the responses section refers to the number of outputs that will be used in the calibration process: in this case, it is just one. The calibration data file has the observational data: in this case, it is a freeform file (e.g. no header or annotation) with ten experiments. For each experiment, there is one standard deviation value indicating the error associated with that experiment.

```
strategy,
        single_method
        tabular_graphics_data

method,
        bayes_calibration queso
#       emulator
#       gp
#        emulator_samples = 50
#       points_file = 'ros.txt' freeform
#        pce
#        sparse_grid_level = 3
          samples = 5000 #seed = 348
          rejection delayed
          metropolis adaptive
          likelihood_scale = 100.0
          output verbose
          #proposal_covariance_scale = 0.1 0.3

variables,
       continuous_design = 2
         lower_bounds = 0. 0.
         upper_bounds = 3. 3.
         initial_point = 1. 1.

interface,
         direct
           analysis_driver = 'text_book'

responses,
        calibration_terms = 1
        calibration_data_file = 'dakota_queso.test10.txt'
          freeform
          num_experiments = 1
          num_replicates = 10
          num_std_deviations = 1
        no_gradients
        no_hessians
```

When the input file shown in figUQ18 is run, Dakota will run the MCMC algorithm and generate a posterior sample of θ in accordance with Bayes Theorem eqnBayesThm and the likelihood function eqnLikelihood. The

MCMC sample output is put into a directory called `outputData` in the directory from which Dakota is run. In addition, the MCMC sample chain is written to a file in the run directory called QuesoOutput.txt. The first columns of this file are the sample inputs, the next columns are the responses, and the final column is the log of the likelihood.

We expect to continue development of Bayesian calibration methods, so check for updates to this capability. The QUESO and GPMSA capabilities in Dakota currently rely on the QUESO library developed by The University of Texas at Austin. This integrated capability is still in prototype form and available to close collaborators of the Dakota team. DREAM is distributed with Dakota. If you are interested in these capabilities, contact the Dakota developers at `dakota-developers@development.sandia.gov`.

Related Topics

Related Keywords

- bayes_calibration : Bayesian calibration
- dream : DREAM (DiffeRential Evolution Adaptive Metropolis)
- chains : Number of chains in DREAM
- crossover_chain_pairs : Number of chains used in crossover.
- gr_threshold : Convergence tolerance for the Gelman-Rubin statistic
- jump_step : Number of generations a long jump step is taken
- num_cr : Number of candidate points for each crossover.
- gpmsa : Gaussian Process Models for Simulation Analysis (GPMSA) Markov Chain Monte Carlo algorithm with Gaussian Process Surrogate
- proposal_covariance_scale : Provide values for the diagonals of the proposal covariance matrix.
- queso : Markov Chain Monte Carlo algorithms from the QUESO package
- mcmc_type : Type of MCMC algorithm used (dram or multilevel).
- dram : MCMC type
- adaptive : Type of Metropolis algorithm
- multilevel : Type of MCMC algorithm
- proposal_covariance_scale : Provide values for the diagonals of the proposal covariance matrix.

4.7.12 nonlinear_least_squares

Description

Dakota's least squares branch currently contains three methods for solving nonlinear least squares problems:

- NL2SOL, a trust-region method that adaptively chooses between two Hessian approximations (Gauss--Newton and Gauss-Newton plus a quasi-Newton approximation to the rest of the Hessian)
- NLSSOL, a sequential quadratic programming (SQP) approach that is from the same algorithm family as NPSOL

- Gauss-Newton, which supplies the Gauss-Newton Hessian approximation to the full-Newton optimizers from OPT++.

The important difference of these algorithms from general-purpose optimization methods is that the response set is defined by calibration terms (e.g. separate terms for each residual), rather than an objective function. Thus, a finer granularity of data is used by least squares solvers as compared to that used by optimizers. This allows the exploitation of the special structure provided by a sum of squares objective function.

Related Topics

Related Keywords

- nl2sol : Trust-region method for nonlinear least squares
- nlssol_sqp : Sequential Quadratic Program for nonlinear least squares

4.7.13 advanced_optimization

Description

empty

Related Topics

- scaling
- multiobjective_methods
- surrogate_based_optimization_methods

Related Keywords

scaling

Description

empty

Related Topics

Related Keywords

multiobjective_methods

Description

empty

Related Topics

Related Keywords

surrogate_based_optimization_methods

Description

empty

Related Topics

Related Keywords

- efficient_global : Global Surrogate Based Optimization, a.k.a. EGO
- surrogate_based_global : Global Surrogate Based Optimization
- surrogate_based_local : Local Surrogate Based Optimization

4.8 advanced_topics

Description

Advanced Dakota capabilities

Related Topics

- advanced_strategies
- advanced_model_recursion
- advanced_simulation_interfaces
- advanced_optimization

Related Keywords

4.8.1 advanced_strategies

Description

empty

Related Topics

Related Keywords

4.8.2 advanced_model_recursion

Description

empty

Related Topics

- hybrid_and_recursions_logic

Related Keywords

hybrid_and_recursions_logic

Description

empty

Related Topics

Related Keywords

4.8.3 advanced_simulation_interfaces

Description

empty

Related Topics

- simulation_failure
- concurrency_and_parallelism

Related Keywords

simulation_failure

Description

empty

Related Topics

Related Keywords

concurrency_and_parallelism

Description

empty

Related Topics

Related Keywords

- processors_per_analysis : Specify the number of processors per analysis when Dakota is run in parallel
- analysis_scheduling : Specify the scheduling of concurrent analyses when Dakota is run in parallel
- master : Specify a dedicated master partition for parallel analysis scheduling

- peer : Specify a peer partition for parallel analysis scheduling
- analysis_servers : Specify the number of analysis servers when Dakota is run in parallel
- asynchronous : Specify analysis driver concurrency, when Dakota is run in serial
- analysis_concurrency : Limit the number of analysis drivers within an evaluation that Dakota will schedule
- evaluation_concurrency : Determine how many concurrent evaluations Dakota will schedule
- local_evaluation_scheduling : Control how local asynchronous jobs are scheduled
- master : Specify a dedicated master partition for parallel evaluation scheduling
- peer : Specify a peer partition for parallel evaluation scheduling
- dynamic : Specify dynamic scheduling in a peer partition when Dakota is run in parallel.
- static : Specify static scheduling in a peer partition when Dakota is run in parallel.
- evaluation_servers : Specify the number of evaluation servers when Dakota is run in parallel
- processors_per_evaluation : Specify the number of processors per evaluation server when Dakota is run in parallel
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- master : Specify a dedicated master partition for parallel iterator scheduling
- peer : Specify a peer partition for parallel iterator scheduling
- iterator_servers : Specify the number of iterator servers when Dakota is run in parallel
- processors_per_iterator : Specify the number of processors per iterator server when Dakota is run in parallel
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- master : Specify a dedicated master partition for parallel iterator scheduling
- peer : Specify a peer partition for parallel iterator scheduling
- iterator_servers : Specify the number of iterator servers when Dakota is run in parallel
- processors_per_iterator : Specify the number of processors per iterator server when Dakota is run in parallel
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- master : Specify a dedicated master partition for parallel iterator scheduling
- peer : Specify a peer partition for parallel iterator scheduling
- iterator_servers : Specify the number of iterator servers when Dakota is run in parallel
- processors_per_iterator : Specify the number of processors per iterator server when Dakota is run in parallel
- iterator_scheduling : Specify the scheduling of concurrent iterators when Dakota is run in parallel
- master : Specify a dedicated master partition for parallel iterator scheduling
- peer : Specify a peer partition for parallel iterator scheduling
- iterator_servers : Specify the number of iterator servers when Dakota is run in parallel
- processors_per_iterator : Specify the number of processors per iterator server when Dakota is run in parallel

4.8.4 advanced_optimization

Description

empty

Related Topics

- scaling
- multiobjective_methods
- surrogate_based_optimization_methods

Related Keywords

scaling

Description

empty

Related Topics

Related Keywords

multiobjective_methods

Description

empty

Related Topics

Related Keywords

surrogate_based_optimization_methods

Description

empty

Related Topics

Related Keywords

- efficient_global : Global Surrogate Based Optimization, a.k.a. EGO
- surrogate_based_global : Global Surrogate Based Optimization
- surrogate_based_local : Local Surrogate Based Optimization

4.9 packages

Description

This topic organizes information about the different software packages (libraries) that are integrated into Dakota

Related Topics

- package_coliny
- package_conmin
- package_ddace
- package_dot
- package_fsudace
- package_hopspack
- package_jega
- package_nlpql
- package_npsol
- package_optpp
- package_psuade
- package_queso
- package_scolib

Related Keywords

4.9.1 package_coliny

Description

SCOLIB (formerly known as COLINY) is a collection of nongradient-based optimizers that support the Common Optimization Library INterface (COLIN). SCOLIB optimizers currently include `coliny_cobyla`, `coliny_direct`, `coliny_ea`, `coliny_pattern_search` and `coliny_solis_wets`. (Yes, the input spec still has "coliny" prepended to the method name.) Additional SCOLIB information is available from https://software.sandia.gov/trac/acro.

SCOLIB solvers now support bound constraints and general nonlinear constraints. Supported nonlinear constraints include both equality and two-sided inequality constraints. SCOLIB solvers do not yet support linear constraints. Most SCOLIB optimizers treat constraints with a simple penalty scheme that adds `constraint_penalty` times the sum of squares of the constraint violations to the objective function. Specific exceptions to this method for handling constraint violations are noted below. (The default value of `constraint_penalty` is 1000.0, except for methods that dynamically adapt their constraint penalty, for which the default value is 1.0.)

The method independent controls for `max_iterations` and `max_function_evaluations` limit the number of major iterations and the number of function evaluations that can be performed during a SCOLIB optimization, respectively. The `convergence_tolerance` control defines the threshold value on relative change in the objective function that indicates convergence. The `output` verbosity specification controls the amount

of information generated by SCOLIB: the `silent`, `quiet`, and `normal` settings correspond to minimal reporting from SCOLIB, whereas the `verbose` setting corresponds to a higher level of information, and `debug` outputs method initialization and a variety of internal SCOLIB diagnostics. The majority of SCOLIB's methods perform independent function evaluations that can directly take advantage of Dakota's parallel capabilities. Only `coliny_solis_wets`, `coliny_cobyla`, and certain configurations of `coliny_pattern_search` are inherently serial. The parallel methods automatically utilize parallel logic when the Dakota configuration supports parallelism. Lastly, neither `speculative` gradients nor linear constraints are currently supported with SCOLIB.

Some SCOLIB methods exploit parallelism through the use of Dakota's concurrent function evaluations. The nature of the algorithms, however, limits the amount of concurrency that can be exploited. The maximum amount of evaluation concurrency that can be leveraged by the various methods is as follows:

- COBYLA: one
- DIRECT: twice the number of variables
- Evolutionary Algorithms: size of the population
- Pattern Search: size of the search pattern
- Solis-Wets: one

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

Each of the SCOLIB methods supports the `solution_target` control, which defines a convergence criterion in which the optimizer will terminate if it finds an objective function value lower than the specified target.

Related Topics

Related Keywords

- coliny_beta : What does this do??????????????
- coliny_cobyla : Constrained Optimization BY Linear Approximations (COBYLA)
- coliny_direct : DIviding RECTangles method
- coliny_ea : Evolutionary Algorithm
- coliny_pattern_search : Pattern search, derivative free optimization method
- coliny_solis_wets : Simple greedy local search method

4.9.2 package_conmin

Description

The CONMIN library [[84] "Vanderplaats, 1973"] is a public domain library of nonlinear programming optimizers, specifically the Fletcher-Reeves conjugate gradient (Dakota's `conmin_frcg` method) method for unconstrained optimization, and the method of feasible directions (Dakota's `conmin_mfd` method) for constrained optimization. As CONMIN was a predecessor to the DOT commercial library, the algorithm controls are very similar.

Related Topics

Related Keywords

- conmin : Access to methods in the CONMIN library
- frcg : A conjugate gradient optimization method
- mfd : Method of feasible directions

4.9.3 package_ddace

Description

The Distributed Design and Analysis of Computer Experiments (DDACE) library provides the following DACE techniques: grid sampling (`grid`), pure random sampling (`random`), orthogonal array sampling (`oas`), latin hypercube sampling (`lhs`), orthogonal array latin hypercube sampling (`oa_lhs`), Box-Behnken (`box_behnken`), and central composite design (`central_composite`).

It is worth noting that there is some overlap in sampling techniques with those available from the nondeterministic branch. The current distinction is that the nondeterministic branch methods are designed to sample within a variety of probability distributions for uncertain variables, whereas the design of experiments methods treat all variables as having uniform distributions. As such, the design of experiments methods are well-suited for performing parametric studies and for generating data sets used in building global approximations, but are not currently suited for assessing the effect of uncertainties characterized with probability distribution. If a design of experiments over both design/state variables (treated as uniform) and uncertain variables (with probability distributions) is desired, then `sampling` can support this with `active all` specified in the Variables specification block.

Related Topics

Related Keywords

- dace : Design and Analysis of Computer Experiments

4.9.4 package_dot

Description

The DOT library [[86] "Vanderplaats Research and Development, 1995"] contains nonlinear programming optimizers, specifically the Broyden-Fletcher-Goldfarb-Shanno (Dakota's `dot_bfgs` method) and Fletcher-Reeves conjugate gradient (Dakota's `dot_frcg` method) methods for unconstrained optimization, and the modified method of feasible directions (Dakota's `dot_mmfd` method), sequential linear programming (Dakota's `dot_slp`

method), and sequential quadratic programming (Dakota's `dot_sqp` method) methods for constrained optimization.

Related Topics

Related Keywords

- dot : Access to methods in the DOT package
- bfgs : A conjugate gradient optimization method
- frcg : A conjugate gradient optimization method
- mmfd : Method of feasible directions
- slp : Sequential Linear Programming
- sqp : Sequential Quadratic Program

4.9.5 package_fsudace

Description

The Florida State University Design and Analysis of Computer Experiments (FSUDace) library provides the following DACE techniques: quasi-Monte Carlo sampling (`fsu_quasi_mc`) based on the Halton sequence (`halton`) or the Hammersley sequence (`hammersley`), and Centroidal Voronoi Tessellation (`fsu_cvt`).

Related Topics

Related Keywords

- fsu_cvt : Design of Computer Experiments - Centroidal Voronoi Tessellation
- fsu_quasi_mc : Design of Computer Experiments - Quasi-Monte Carlo sampling

4.9.6 package_hopspack

Description

The HOPSPACK software [[70] "Plantenga, 2009"] contains the asynchronous parallel pattern search (APPS) algorithm [[37] "Gray and Kolda, 2006"]. It can handle unconstrained problems as well as those with bound constraints, linear constraints, and general nonlinear constraints.

HOPSPACK is available to the public under the GNU LGPL and the source code is included with Dakota. HOPSPACK-specific software documentation is available from `https://software.sandia.gov/trac/hopspack`.

Related Topics

Related Keywords

- asynch_pattern_search : Pattern search, derivative free optimization method

4.9.7 package_jega

Description

The JEGA library [[18] "Eddy and Lewis, 2001"] contains two global optimization methods. The first is a Multi-objective Genetic Algorithm (MOGA) which performs Pareto optimization. The second is a Single-objective Genetic Algorithm (SOGA) which performs optimization on a single objective function. Both methods support general constraints and a mixture of real and discrete variables. The JEGA library was written by John Eddy, currently a member of the technical staff in the System Readiness and Sustainment Technologies department at Sandia National Laboratories in Albuquerque. These algorithms are accessed as `moga` and `soga` within Dakota.

Related Topics

Related Keywords

- moga : Multi-objective Genetic Algorithm (a.k.a Evolutionary Algorithm)
- soga : Single-objective Genetic Algorithm (a.k.a Evolutionary Algorithm)

4.9.8 package_nlpql

Description

The NLPQL library is a commercially-licensed library containing a sequential quadratic programming (SQP) optimizer, specified as Dakota's `nlpql_sqp` method, for constrained optimization. The particular implementation used is NLPQLP [[75] "Schittkowski, 2004"], a variant with distributed and non-monotone line search.

Related Topics

Related Keywords

- nlpql_sqp : Sequential Quadratic Program

4.9.9 package_npsol

Description

The NPSOL library [[33] "Gill et al., 1986"] contains a sequential quadratic programming (SQP) implementation (the `npsol_sqp` method). SQP is a nonlinear programming optimizer for constrained minimization.

Related Topics

Related Keywords

- npsol_sqp : Sequential Quadratic Program

4.9.10 package_optpp

Description

The OPT++ library [[61] "Meza et al., 2007"] contains primarily gradient-based nonlinear programming optimizers for unconstrained, bound-constrained, and nonlinearly constrained minimization: Polak-Ribiere conjugate

gradient (Dakota's optpp_cg method), quasi-Newton (Dakota's optpp_q_newton method), finite difference Newton (Dakota's optpp_fd_newton method), and full Newton (Dakota's optpp_newton method).

The conjugate gradient method is strictly unconstrained, and each of the Newton-based methods are automatically bound to the appropriate OPT++ algorithm based on the user constraint specification (unconstrained, bound-constrained, or generally-constrained). In the generally-constrained case, the Newton methods use a nonlinear interior-point approach to manage the constraints. The library also contains a direct search algorithm, PDS (parallel direct search, Dakota's optpp_pds method), which supports bound constraints.

Method Independent Controls

These are specified directly under the method block.

1. max_iterations
2. max_function_evaluations
3. convergence_tolerance
4. output
5. speculative

Concurrency

OPT++'s gradient-based methods are not parallel algorithms and cannot directly take advantage of concurrent function evaluations. However, if numerical_gradients with method_source dakota is specified, a parallel Dakota configuration can utilize concurrent evaluations for the finite difference gradient computations.

Constraints

Linear constraint specifications are supported by each of the Newton methods (optpp_newton, optpp_q_newton, optpp_fd_newton, and optpp_g_newton)

optpp_cg must be unconstrained

optpp_pds can be, at most, bound-constrained.

Related Topics

Related Keywords

- optpp_cg : A conjugate gradient optimization method
- optpp_fd_newton : Finite Difference Newton optimization method
- optpp_g_newton : Newton method based least-squares calbration
- optpp_newton : Newton method based optimization
- optpp_pds : Simplex-based derivative free optimization method
- optpp_q_newton : Quasi-Newton optimization method

4.9.11 package_psuade

Description

The Problem Solving Environment for Uncertainty Analysis and Design Exploration (PSUADE) is a Lawrence Livermore National Laboratory tool for metamodeling, sensitivity analysis, uncertainty quantification, and optimization. Its features include non-intrusive and parallel function evaluations, sampling and analysis methods, an integrated design and analysis framework, global optimization, numerical integration, response surfaces (MARS and higher order regressions), graphical output with Pgplot or Matlab, and fault tolerance [[82] "C.H. Tong, 2005"].

Related Topics

Related Keywords

- psuade_moat : Morris One-at-a-Time

4.9.12 package_queso

Description

QUESO stands for Quantification of Uncertainty for Estimation, Simulation, and Optimization. It supports Bayesian calibration methods. It is developed at The University of Texas at Austin.

Related Topics

Related Keywords

- bayes_calibration : Bayesian calibration
- gpmsa : Gaussian Process Models for Simulation Analysis (GPMSA) Markov Chain Monte Carlo algorithm with Gaussian Process Surrogate
- queso : Markov Chain Monte Carlo algorithms from the QUESO package

4.9.13 package_scolib

Description

SCOLIB (formerly known as COLINY) is a collection of nongradient-based optimizers that support the Common Optimization Library INterface (COLIN). SCOLIB optimizers currently include `coliny_cobyla`, `coliny_direct`, `coliny_ea`, `coliny_pattern_search` and `coliny_solis_wets`. (Yes, the input spec still has "coliny" prepended to the method name.) Additional SCOLIB information is available from https://software.sandia.gov/trac/acro.

SCOLIB solvers now support bound constraints and general nonlinear constraints. Supported nonlinear constraints include both equality and two-sided inequality constraints. SCOLIB solvers do not yet support linear constraints. Most SCOLIB optimizers treat constraints with a simple penalty scheme that adds `constraint_penalty` times the sum of squares of the constraint violations to the objective function. Specific exceptions to this method for handling constraint violations are noted below. (The default value of `constraint_penalty` is 1000.0, except for methods that dynamically adapt their constraint penalty, for which the default value is 1.0.)

The method independent controls for `max_iterations` and `max_function_evaluations` limit the number of major iterations and the number of function evaluations that can be performed during a SCOLIB optimization, respectively. The `convergence_tolerance` control defines the threshold value on relative change in the objective function that indicates convergence. The `output` verbosity specification controls the amount of information generated by SCOLIB: the `silent`, `quiet`, and `normal` settings correspond to minimal reporting from SCOLIB, whereas the `verbose` setting corresponds to a higher level of information, and `debug` outputs method initialization and a variety of internal SCOLIB diagnostics. The majority of SCOLIB's methods perform independent function evaluations that can directly take advantage of Dakota's parallel capabilities. Only `coliny_solis_wets`, `coliny_cobyla`, and certain configurations of `coliny_pattern_search` are inherently serial. The parallel methods automatically utilize parallel logic when the Dakota configuration supports parallelism. Lastly, neither `speculative` gradients nor linear constraints are currently supported with SCOLIB.

Some SCOLIB methods exploit parallelism through the use of Dakota's concurrent function evaluations. The nature of the algorithms, however, limits the amount of concurrency that can be exploited. The maximum amount of evaluation concurrency that can be leveraged by the various methods is as follows:

- COBYLA: one
- DIRECT: twice the number of variables
- Evolutionary Algorithms: size of the population
- Pattern Search: size of the search pattern
- Solis-Wets: one

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

Each of the SCOLIB methods supports the `solution_target` control, which defines a convergence criterion in which the optimizer will terminate if it finds an objective function value lower than the specified target.

Related Topics

Related Keywords

- coliny_beta : What does this do????????????????
- coliny_cobyla : Constrained Optimization BY Linear Approximations (COBYLA)
- coliny_direct : DIviding RECTangles method
- coliny_ea : Evolutionary Algorithm
- coliny_pattern_search : Pattern search, derivative free optimization method
- coliny_solis_wets : Simple greedy local search method

Chapter 5

Keywords Area

This page lists the six blocks. From here, you can navigate to every keyword.

- environment
- method
- model
- variables
- interface
- responses

Introduction to Dakota Keywords

In Dakota, the *environment* manages execution modes and I/O streams and defines the top-level iterator. Generally speaking, an iterator contains a model and a model contains a set of *variables*, an *interface*, and a set of *responses*. An iterator repeatedly operates on the model to map the variables into responses using the interface. Each of these six components (environment, method, model, variables, interface, and responses) are separate specifications in the user's input file, and as a whole, determine the study to be performed during an execution of the Dakota software.

A Dakota execution is limited to a single environment, but may involve multiple methods and multiple models. In particular, advanced iterators (i.e., meta- and component-based iterators) and advanced models (i.e., nested and surrogate models) may specialize to include recursions with additional sub-iterators and sub-models. Since each model may contain its own variables, interface, and responses, there may be multiple specifications of the method, model, variables, interface, and responses sections.

Keyword Pages

Every Dakota keyword has its own page in this manual. The page describes:

- Whether the keyword takes `ARGUMENTS`, and the data type Additional notes about `ARGUMENTS` can be found here: Specifying Arguments.
- Whether it has an `ALIAS`
- Which additional keywords can be specified to change its behavior
- Which of these additional keywords are required or optional
- Additional information about how to use the keyword in an input file

5.1 environment

- Keywords Area
- environment

Top-level settings for Dakota execution

Topics

This keyword is related to the topics:

- block

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		graphics	Display a 2D graphics window of variables and responses
	Optional		tabular_graphics_-data	Write a tabular results file with variable and response history
	Optional		output_precision	Control the output precision
	Optional		results_output	(Experimental) Write a summary file containing the final results
	Optional		top_method_-pointer	Identify which method leads the Dakota study

Description

The environment section in a Dakota input file is optional. It specifies the top-level solution environment, optionally including output controls and identification of the primary iterative method. The output-related keywords address graphics, generation of tabular and results data, and the precision of numerical output.

The `top_method_pointer` keyword is typically used in Dakota studies consisting of more than one method block to clearly indicate which is the leading method. This method provides the starting point for the iteration. The corresponding method specification may recurse with additional sub-method pointers in the case of "meta-iteration" (see method) or may specify a single method without recursion. Either case will ultimately result in identification of one or more model specifications using `model_pointer`, which again may or may not involve

further recursion (see nested and surrogate for recursion cases). Each of the model specifications identify the variables and responses specifications (using variables_pointer and responses_pointer) that are used to build the model, and depending on the type of model, may also identify an interface specification (for example, using interface_pointer). If one of these specifications does not provide an optional pointer, then that component will be constructed using the last specification parsed.

When the environment block is omitted, the top level method will be inferred as follows: When a single method is specified, there is no ambiguity and the sole method will be the top method. When multiple methods are specified, the top level method will be deduced from the hierarchical relationships implied by method pointers. If this inference is not well defined (e.g., there are multiple method specifications without any pointer relationship), then the default behavior is to employ the last method specification parsed.

5.1.1 graphics

- Keywords Area
- environment
- graphics

Display a 2D graphics window of variables and responses

Topics

This keyword is related to the topics:

- dakota_output

Specification

Alias: none

Argument(s): none

Description

For most studies, the `graphics` flag activates a 2D graphics window containing history plots for the variables and response functions in the study. This window is updated in an event loop with approximately a 2 second cycle time. Some study types such as surrogate-based optimization or local reliability specialize the use of the graphics window.

There is no dependence between the `graphics` flag and the `tabular_graphics_data` flag; they may be used independently or concurrently.

See Also

These keywords may also be of interest:

- tabular_graphics_data

5.1.2 tabular_graphics_data

- Keywords Area
- environment
- tabular_graphics_data

Write a tabular results file with variable and response history

Topics

This keyword is related to the topics:

- dakota_output

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		tabular_graphics_-file	The name of the tabular data output file

Description

Specifying the `tabular_graphics_data` flag writes to a data file the same variable and response function history data plotted when using the `graphics` flag. Within the generated data file, the variables and response functions appear as columns and each function evaluation provides a new table row. This capability is most useful for post-processing of Dakota results with third-party graphics tools such as MatLab, Excel, Tecplot, etc.

There is no dependence between the `graphics` flag and the `tabular_graphics_data` flag; they may be used independently or concurrently.

See Also

These keywords may also be of interest:

- graphics

tabular_graphics_file

- Keywords Area
- environment
- tabular_graphics_data
- tabular_graphics_file

The name of the tabular data output file

Topics

This keyword is related to the topics:

- dakota_output

Specification

Alias: none
Argument(s): STRING

Description

The optional `tabular_graphics_file` keyword specifies a name to use for the tabular data file, overriding the default `dakota_tabular.dat`.

5.1.3 output_precision

- Keywords Area
- environment
- output_precision

Control the output precision

Topics

This keyword is related to the topics:

- dakota_output

Specification

Alias: none
Argument(s): INTEGER

Description

The precision of numeric output precision can be set dwith `output_precision`, with an upper limit of 16. When not specified, most Dakota output will default to a precision of 10, though filesystem interfaces and pre-run output use higher precision for accuracy and better results reproducibility.

5.1.4 results_output

- Keywords Area
- environment
- results_output

(Experimental) Write a summary file containing the final results

Topics

This keyword is related to the topics:

- dakota_output

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		results_output_file	The base file name of the results file

Description

Final results from a Dakota study can be output to dakota_results.txt by specifying `results_output` (optionally specifying an alternate file name with `results_output_filename`). The current experimental text file format is hierarchical and a precursor to planned output to structured text formats such as XML or YAML, and binary formats such as HDF5. The contents, organization, and format of results files are all under active development and are subject to change.

results_output_file

- Keywords Area
- environment
- results_output
- results_output_file

The base file name of the results file

Topics

This keyword is related to the topics:

- dakota_output

Specification

Alias: none
Argument(s): STRING

Description

Default file name is `dakota_results.txt`

5.1.5 top_method_pointer

- Keywords Area
- environment
- top_method_pointer

Identify which method leads the Dakota study

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: method_pointer
Argument(s): STRING

Description

An optional `top_method_pointer` specification may be used to point to a particular method specification that will lead the Dakota analysis. The associated string must be a method identifier specified via id_method. If `top_method_pointer` is not used, then it will be inferred as decribed in environment (no `top_method_pointer` within an environment specification is treated the same as no environment specification).

See Also

These keywords may also be of interest:

- id_method

5.2 method

- Keywords Area
- method

Specifies the name and behavior of a method

Topics

This keyword is related to the topics:

- block
- not_yet_reviewed

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		id_method	Give the block a name, in case of multiple method blocks
	Optional		output	Control how much method information is written to the screen and output file
	Optional		max_iterations	Stopping criteria based on number of iterations
	Optional		max_function_-evaluations	Stopping criteria based on number of function evaluations
	Optional		speculative	Compute speculative gradients
	Optional		convergence_-tolerance	Stopping criterion based on convergence of the objective function
	Optional		constraint_-tolerance	The maximum allowable value of constraint violation still considered to be feasible
	Optional		scaling	Turn on scaling for variables, responses, and constraints
	Optional		final_solutions	Number of designs returned as the best solutions
			hybrid	Strategy in which a set of methods synergistically seek an optimal design

			multi_start	Multi-Start Optimization Method
			pareto_set	Pareto set optimization
			surrogate_based_-local	Local Surrogate Based Optimization
			surrogate_based_-global	Global Surrogate Based Optimization
			dot_frcg	A conjugate gradient optimization method
			dot_mmfd	Method of feasible directions
			dot_bfgs	A conjugate gradient optimization method
			dot_slp	Sequential Linear Programming
			dot_sqp	Sequential Quadratic Program
			dot	Access to methods in the DOT package
			conmin_frcg	A conjugate gradient optimization method
			conmin_mfd	Method of feasible directions
			conmin	Access to methods in the CONMIN library
			dl_solver	There is ZERO mention of this in the reference manual??

npsol_sqp	Sequential Quadratic Program
nlssol_sqp	Sequential Quadratic Program for nonlinear least squares
stanford	Select methods from the Stanford package
nlpql_sqp	Sequential Quadratic Program
optpp_cg	A conjugate gradient optimization method
optpp_q_newton	Quasi-Newton optimization method
optpp_fd_newton	Finite Difference Newton optimization method
optpp_g_newton	Newton method based least-squares calbration
optpp_newton	Newton method based optimization
optpp_pds	Simplex-based derivative free optimization method
asynch_pattern_-search	Pattern search, derivative free optimization method
mesh_adaptive_-search	Mesh Adaptive Direct Search Algorithm
moga	Multi-objective Genetic Algorithm (a.k.a Evolutionary Algorithm)

			soga	Single-objective Genetic Algorithm (a.k.a Evolutionary Algorithm)
			coliny_pattern_-search	Pattern search, derivative free optimization method
			coliny_solis_wets	Simple greedy local search method
			coliny_cobyla	Constrained Optimization BY Linear Approximations (COBYLA)
			coliny_direct	DIviding RECTangles method
			coliny_ea	Evolutionary Algorithm
			coliny_beta	What does this do??????????????
			nl2sol	Trust-region method for nonlinear least squares
			nonlinear_cg	This is not mentioned in the reference manual
			ncsu_direct	DIviding RECTangles method
			genie_opt_darts	
			genie_direct	
			efficient_global	Global Surrogate Based Optimization, a.k.a. EGO
			polynomial_chaos	Uncertainty quantification using polynomial chaos expansions

				stoch_collocation	Uncertainty quantification with stochastic collocation
				sampling	Uncertainty quantification based on random sampling of the variables
				importance_-sampling	Importance sampling
				gpais	Gaussian Process Adaptive Importance Sampling
				adaptive_sampling	Build a GP surrogate and refine it adaptively
				pof_darts	Probability-of--Failure (POF) darts is a novel method for estimating the probability of failure based on random sphere-packing.
				efficient_subspace	This is not mentioned in the reference manual
				global_evidence	Evidence theory with evidence measures computed with global optimization methods
				global_interval_est	Interval analysis using global optimization methods
				bayes_calibration	Bayesian calibration
				dace	Design and Analysis of Computer Experiments

			fsu_cvt	Design of Computer Experiments - Centroidal Voronoi Tessellation
			psuade_moat	Morris One-at-a-Time
			local_evidence	Evidence theory with evidence measures computed with local optimization methods
			local_interval_est	Interval analysis using local optimization
			local_reliability	Local reliability method
			global_reliability	Global reliability methods
			fsu_quasi_mc	Design of Computer Experiments - Quasi-Monte Carlo sampling
			vector_parameter_-study	Parameter study - samples points along a vector
			list_parameter_-study	Parameter study specified by a list of variable values
			centered_-parameter_study	Parameter study - points moving out from a center point
			multidim_-parameter_study	Parameter study - make a full factorial grid of study points
			richardson_extrap	Estimate order of convergence of a response as model fidelity increases

Description

The method block(s) in a Dakota input file each specify the various keywords to select a method and control its behavior. The terms ”method” and ”iterator” can be used interchangeably.

The Methods

Each `method` block must select one method.

Starting with Dakota v6.0, the methods are grouped into two types: standard methods and multi-component

methods.

The standard methods are stand-alone and self-contained in the sense that they only require a model to perform a study. They do not call other methods. While methods such as `polynomial_chaos` and `efficient_global` internally utilize multiple iterator and surrogate model components, these components are generally hidden from user control due to restrictions on modularity; thus, these methods are stand-alone.

The multi-component group of methods provides a higher level "meta-algorithm" that points to other methods and models that support sub-iteration. For example, in a sequential hybrid method, the `hybrid` method specification must identify a list of subordinate methods, and the "meta-algorithm" executes these methods in sequence and transfers information between them. Surrogate-based minimizers provide another example in that they point both to other methods (e.g. what optimization method is used to solve the approximate subproblem) as well as to models (e.g. what type of surrogate model is employed). Multi-component methods generally provide some level of "plug and play" modularity, through their flexible support of a variety of method and model selections.

Component-Based Iterator Commands

Component-based iterator specifications include hybrid, multi-start, pareto set, surrogate-based local, surrogate-based global, and branch and bound methods. Whereas a standard iterator specification only needs an optional model pointer string (specified with `model_pointer`), component-based iterator specifications can include method pointer, method name, and model pointer specifications in order to define the components employed in the "meta-iteration." In particular, these specifications identify one or more methods (by pointer or by name) to specify the subordinate iterators that will be used in the top-level algorithm. Identifying a sub-iterator by name instead of by pointer is a lightweight option that relaxes the need for a separate method specification for the sub-iterator; however, a model pointer may be required in this case to provide the specification connectivity normally supported by the method pointer. Refer to these individual method descriptions for specific requirements for these advanced methods.

Method Independent Controls

In addition to the method, there are 10 optional keywords, which are referred to as method independent controls. These controls are valid for enough methods that it was reasonable to pull them out of the method dependent blocks and consolidate the specifications, however, they are NOT universally respected by all methods.

Examples

Several examples follow. The first example shows a minimal specification for an optimization method.

```
method
  dot_sqp
```

This example uses all of the defaults for this method.

A more sophisticated example would be

```
method,
  id_method = 'NLP1'
  dot_sqp
    max_iterations = 50
    convergence_tolerance = 1e-4
    output verbose
    model_pointer = 'M1'
```

This example demonstrates the use of identifiers and pointers as well as some method independent and method dependent controls for the sequential quadratic programming (SQP) algorithm from the DOT library. The `max_iterations`, `convergence_tolerance`, and `output` settings are method independent controls, in that they are defined for a variety of methods (see dot for usage of these controls).

The next example shows a specification for a least squares method.

```
method
  optpp_g_newton
    max_iterations = 10
```

```
    convergence_tolerance = 1.e-8
    search_method trust_region
    gradient_tolerance = 1.e-6
```

Some of the same method independent controls are present along with several method dependent controls (`search_method` and `gradient_tolerance`) which are only meaningful for OPT++ methods (see package_optpp).

The next example shows a specification for a nondeterministic method with several method dependent controls (refer to sampling).

```
method
  sampling
    samples = 100
    seed = 12345
    sample_type lhs
    response_levels = 1000. 500.
```

The last example shows a specification for a parameter study method where, again, each of the controls are method dependent (refer to vector_parameter_study).

```
method
  vector_parameter_study
    step_vector = 1. 1. 1.
    num_steps = 10
```

5.2.1 id_method

- Keywords Area
- method
- id_method

Give the block a name, in case of multiple method blocks

Topics

This keyword is related to the topics:

- block_identifier
- method_independent_controls

Specification

Alias: none

Argument(s): STRING

Description

The method identifier string is supplied with `id_method` and is used to provide a unique identifier string for use with environment or meta-iterator specifications (refer to environment). It is appropriate to omit a method identifier string if only one method is included in the input file, since the single method to use is unambiguous in this case.

5.2.2 output

- Keywords Area
- method
- output

Control how much method information is written to the screen and output file

Topics

This keyword is related to the topics:

- dakota_output
- method_independent_controls

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword debug	Dakota Keyword Description Level 5 of 5 - maximum
	Required*(Choose One)*	**Group 1**	verbose	Level 4 of 5 - more than normal
			normal	Level 3 of 5 - default
			quiet	Level 2 of 5 - less than normal
			silent	Level 1 of 5 - minimum

Description

Choose from a total of five output levels during the course of a Dakota study. If there is no user specification for output verbosity, then the default setting is `normal`.

Specific mappings are as follows:

- `silent` (i.e., really quiet): silent iterators, silent model, silent interface, quiet approximation, quiet file operations
- `quiet`: quiet iterators, quiet model, quiet interface, quiet approximation, quiet file operations
- `normal`: normal iterators, normal model, normal interface, quiet approximation, quiet file operations
- `verbose`: verbose iterators, normal model, verbose interface, verbose approximation, verbose file operations
- `debug` (i.e., really verbose): debug iterators, normal model, debug interface, verbose approximation, verbose file operations

Note that iterators and interfaces utilize the full granularity in verbosity, whereas models, approximations, and file operations do not. With respect to iterator verbosity, different iterators implement this control in slightly different ways (as described below in the method independent controls descriptions for each iterator), however the meaning is consistent.

For models, interfaces, approximations, and file operations, `quiet` suppresses parameter and response set reporting and `silent` further suppresses function evaluation headers and scheduling output. Similarly, `verbose` adds file management, approximation evaluation, and global approximation coefficient details, and `debug` further adds diagnostics from nonblocking schedulers.

debug

- Keywords Area
- method
- output
- debug

Level 5 of 5 - maximum

Specification

Alias: none
Argument(s): none

Description

This is described on output

verbose

- Keywords Area
- method
- output
- verbose

Level 4 of 5 - more than normal

Specification

Alias: none
Argument(s): none

Description

This is described on output

normal

- Keywords Area
- method
- output
- normal

Level 3 of 5 - default

Specification

Alias: none
Argument(s): none

Description

This is described on output

quiet

- Keywords Area
- method
- output
- quiet

Level 2 of 5 - less than normal

Specification

Alias: none
Argument(s): none

Description

This is described on output

silent

- Keywords Area
- method
- output
- silent

Level 1 of 5 - minimum

Specification

Alias: none
Argument(s): none

Description

This is described on output

5.2.3 max_iterations

- Keywords Area
- method
- max_iterations

Stopping criteria based on number of iterations

Topics

This keyword is related to the topics:

- method_independent_controls

Specification

Alias: none
Argument(s): INTEGER

Description

The maximum number of iterations.
The default for `max_iterations` is 100.

See Also

These keywords may also be of interest:

- max_function_evaluations

5.2.4 max_function_evaluations

- Keywords Area
- method
- max_function_evaluations

Stopping criteria based on number of function evaluations

Topics

This keyword is related to the topics:

- method_independent_controls

Specification

Alias: none
Argument(s): INTEGER

Description

The maximum number of function evaluations.
The default for `max_function_evaluations` is 1000.

See Also

These keywords may also be of interest:

- max_iterations

5.2.5 speculative

- Keywords Area
- method
- speculative

Compute speculative gradients

Topics

This keyword is related to the topics:

- method_independent_controls

Specification

Alias: none
Argument(s): none

Description

When performing gradient-based optimization in parallel, `speculative` gradients can be selected to address the load imbalance that can occur between gradient evaluation and line search phases. In a typical gradient-based optimization, the line search phase consists primarily of evaluating the objective function and any constraints at a trial point, and then testing the trial point for a sufficient decrease in the objective function value and/or constraint violation. If a sufficient decrease is not observed, then one or more additional trial points may be attempted sequentially. However, if the trial point is accepted then the line search phase is complete and the gradient evaluation phase begins. By speculating that the gradient information associated with a given line search

trial point will be used later, additional coarse grained parallelism can be introduced by computing the gradient information (either by finite difference or analytically) in parallel, at the same time as the line search phase trial-point function values. This balances the total amount of computation to be performed at each design point and allows for efficient utilization of multiple processors. While the total amount of work performed will generally increase (since some speculative gradients will not be used when a trial point is rejected in the line search phase), the run time will usually decrease (since gradient evaluations needed at the start of each new optimization cycle were already performed in parallel during the line search phase). Refer to [[13] "Byrd et al., 1998"] for additional details. The speculative specification is implemented for the gradient-based optimizers in the DOT, CONMIN, and OPT++ libraries, and it can be used with dakota numerical or analytic gradient selections in the responses specification (refer to responses gradient section for information on these specifications). It should not be selected with vendor numerical gradients since vendor internal finite difference algorithms have not been modified for this purpose. In full-Newton approaches, the Hessian is also computed speculatively. NPSOL and NLSSOL do not support speculative gradients, as their gradient-based line search in user-supplied gradient mode (dakota numerical or analytic gradients) is a superior approach for load-balanced parallel execution.

The `speculative` specification enables speculative computation of gradient and/or Hessian information, where applicable, for parallel optimization studies. By speculating that the derivative information at the current point will be used later, the complete data set (all available gradient/Hessian information) can be computed on every function evaluation. While some of these computations will be wasted, the positive effects are a consistent parallel load balance and usually shorter wall clock time. The `speculative` specification is applicable only when parallelism in the gradient calculations can be exploited by Dakota (it will be ignored for `vendor numerical` gradients).

5.2.6 convergence_tolerance

- Keywords Area
- method
- convergence_tolerance

Stopping criterion based on convergence of the objective function

Topics

This keyword is related to the topics:

- method_independent_controls

Specification

Alias: none
Argument(s): REAL

Description

The `convergence_tolerance` specification provides a real value for controlling the termination of iteration.

It is a **relative convergence tolerance** for the objective function; i.e., if the change in the objective function between successive iterations divided by the previous objective function is less than the amount specified by convergence_tolerance, then this convergence criterion is satisfied on the current iteration.

Therefore, permissible values are between 0 and 1, non-inclusive.

Behavior Varies by Package/Library

This control is used with optimization and least squares iterators (DOT, CONMIN, NPSOL, NLSSOL, OPT++, and SCOLIB) and is not used within the uncertainty quantification, design of experiments, or parameter study iterator branches.

Since no progress may be made on one iteration followed by significant progress on a subsequent iteration, some libraries require that the convergence tolerance be satisfied on two or more consecutive iterations prior to termination of iteration.

Notes on each library:

- DOT: must be satisfied for two consecutive iterations
- NPSOL: defines an internal optimality tolerance which is used in evaluating if an iterate satisfies the first-order Kuhn-Tucker conditions for a minimum. The magnitude of `convergence_tolerance` approximately specifies the number of significant digits of accuracy desired in the final objective function (e.g., `convergence_tolerance = 1.e-6` will result in approximately six digits of accuracy in the final objective function).
- NL2SOL: See nl2sol

5.2.7 constraint_tolerance

- Keywords Area
- method
- constraint_tolerance

The maximum allowable value of constraint violation still considered to be feasible

Topics

This keyword is related to the topics:

- method_independent_controls

Specification

Alias: none

Argument(s): REAL

Description

The `constraint_tolerance` specification determines the maximum allowable value of infeasibility that any constraint in an optimization problem may possess and still be considered to be satisfied.

If a constraint function is greater than this value then it is considered to be violated by the optimization algorithm. This specification gives some control over how tightly the constraints will be satisfied at convergence of the algorithm. However, if the value is set too small the algorithm may terminate with one or more constraints being violated.

This specification is currently meaningful for the NPSOL, NLSSOL, DOT and CONMIN constrained optimizers.

Defaults

Defaults can vary depending on the method.

- DOT constrained optimizers: 0.003

- NPSOL: dependent upon the machine precision, typically on the order of `1.e-8` for double precision computations

5.2.8 scaling

- Keywords Area
- method
- scaling

Turn on scaling for variables, responses, and constraints

Topics

This keyword is related to the topics:

- not_yet_reviewed
- method_independent_controls

Specification

Alias: none
Argument(s): none

Description

Some of the optimization and calibration methods support scaling of continuous design variables, objective functions, calibration terms, and constraints. This is activated by by providing the `scaling` keyword. Discrete variable scaling is not supported.

When scaling is enabled, variables, functions, gradients, Hessians, etc., are transformed such that the method iterates in scaled variable space, whereas evaluations of the computational model as specified in the interface are performed on the original problem scale. Therefore using scaling does not require rewriting the interface to the simulation code.

Scaling also requires the specification of additional keywords which are found in the method, variables, and responses blocks. When the `scaling` keyword is omitted, all `_scale_types` and `*_scales` specifications are ignored in the method, variables, and responses sections.

This page describes the usage of all scaling related keywords. The additional keywords come in pairs, one pair for each set of quantities to be scaled. These quantities can be constraint equations, variables, or responses.

- a `*scales` keyword, which gives characteristic values
- a `*scale_type` keyword, which determines how to use the characteristic values

The pair of keywords both take argument(s), and the length of the arguments can either be zero, one, or equal to the number of quantities to be scaled. If one argument is given, it will apply to all quantities in the set. See the examples below.

Scale Types

There are four scale types:

1. `none` (default) - no scaling, value of `*scales` keyword is ignored

2. `value` - multiplicative scaling

3. `auto` - automatic scaling

   First the quantity is scaled by the characteristic value, then automatic scaling will be attempted according to the following scheme:

   - two-sided bounds scaled into the interval [0,1];
   - one-sided bound or targets are scaled by the characteristic value, moving the bound or target to 1 and changing the sense of inequalities where necessary;
   - no bounds or targets: no automatic scaling possible, therefore no scaling for this component

   Automatic scaling is not available for objective functions nor calibration terms since they lack bound constraints. Futher, when automatically scaled, linear constraints are scaled by characteristic values only, not affinely scaled into [0,1].

4. `log` - logarithmic scaling

   First, any characteristic values from the optional `*_scales` specification are applied. Then logarithm base 10 scaling is applied.

   Logarithmic scaling is not available for linear constraints.

   When continuous design variables are log scaled, linear constraints are not allowed.

Scales

The `*scales` keywords are used to specify the characteristic values. These must be non-zero real numbers. The numbers are used according to the corresponding `*scale_type`, as described above.

Depending on the scale type, the characteristic values may be required or optional.

- `none`, `auto`, `log` - optional
- `value` - required.

A warning is issued if scaling would result in division by a value smaller in magnitude than `1.0e10*DBL_MIN`. User-provided values violating this lower bound are accepted unaltered, whereas for automatically calculated scaling, the lower bound is enforced.

Examples

The two examples below are equivalent:

```
responses
  objective_functions 3
  sense "maximize"
  primary_scale_types = "value"
  primary_scales = 1 1 100

responses
  objective_functions 3
  sense "maximize"
  primary_scale_types = "value" "value" "value"
  primary_scales = 1 1 100
```

5.2.9 final_solutions

- Keywords Area
- method
- final_solutions

Number of designs returned as the best solutions

Topics

This keyword is related to the topics:

- method_independent_controls

Specification

Alias: none
Argument(s): INTEGER

Description

The `final_solutions` controls the number of final solutions returned by the iterator as the best solutions.

For most optimizers, this is one, but some optimizers can produce multiple solutions (e.g. genetic algorithms).

When using a `hybrid` strategy, the number of final solutions dictates how many solutions are passed from one method to another.

Examples

In the case of sampling methods, if one specifies 100 samples (for example) but also specifies `final_solutions` = 5, the five best solutions (in order of lowest response function value) are returned.

5.2.10 hybrid

- Keywords Area
- method
- hybrid

Strategy in which a set of methods synergistically seek an optimal design

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Required*(Choose One)*	**Group 1**	sequential	Methods are run one at a time, in sequence
			embedded	A subordinate local method provides periodic refinements to a top-level global method
			collaborative	Multiple methods run concurrently and share information
	Optional		iterator_servers	Specify the number of iterator servers when Dakota is run in parallel
	Optional		iterator_scheduling	Specify the scheduling of concurrent iterators when Dakota is run in parallel
	Optional		processors_per_-iterator	Specify the number of processors per iterator server when Dakota is run in parallel

Description

In a hybrid minimization method (`hybrid`), a set of methods synergistically seek an optimal design. The relationships among the methods are categorized as:

- collaborative
- embedded
- sequential

The goal in each case is to exploit the strengths of different optimization and nonlinear least squares algorithms at different stages of the minimization process. Global + local hybrids (e.g., genetic algorithms combined with nonlinear programming) are a common example in which the desire for identification of a global optimum is balanced with the need for efficient navigation to a local optimum.

sequential

- Keywords Area
- method

- hybrid
- sequential

Methods are run one at a time, in sequence

Specification

Alias: uncoupled

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	method_name_list	
			method_pointer_list	

Description

In the `sequential` approach, methods are run one at a time, in sequence. The best solutions from one method are used to initialize the next method.

A list of method strings supplied with the `method_list` specification specifies the identity and sequence of iterators to be used. Any number of iterators may be specified.

Method switching is managed through the separate convergence controls of each method. The number of solutions transferred between methods is specified by the particular method through its `final_solutions` method control.

For example, if one sets up a two-level study with a first method that generates multiple solutions such as a genetic algorithm, followed by a second method that is initialized only at a single point such as a gradient-based algorithm, it is possible to take the multiple solutions generated by the first method and create several instances of the second method, each one with a different initial starting point.

The logic governing the transfer of multiple solutions between methods is as follows:

- if one solution is returned from method A, then one solution is transferred to method B.
- If multiple solutions are returned from method A, and method B can accept multiple solutions as input (for example, as a genetic algorithm population), then one instance of method B is initialized with multiple solutions.
- If multiple solutions are returned from method A but method B only can accept one initial starting point, then method B is run several times, each one with a separate starting point from the results of method A.

method_name_list

- Keywords Area
- method
- hybrid
- sequential
- method_name_list

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- block_pointer
- not_yet_reviewed

Specification

Alias: none
Argument(s): STRINGLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		model_pointer_list	

model_pointer_list

- Keywords Area
- method
- hybrid
- sequential
- method_name_list
- model_pointer_list

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- block_pointer
- not_yet_reviewed

Specification

Alias: none
Argument(s): STRING

method_pointer_list

- Keywords Area
- method
- hybrid
- sequential
- method_pointer_list

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- block_pointer
- not_yet_reviewed

Specification

Alias: none
Argument(s): STRINGLIST

embedded

- Keywords Area
- method
- hybrid
- embedded

A subordinate local method provides periodic refinements to a top-level global method

Topics

This keyword is related to the topics:

- problem

Specification

Alias: coupled
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required(*Choose One*)	**Group 1**	global_method_-name	Specify the global method
			global_method_-pointer	Specify the global method
	Required(*Choose One*)	**Group 2**	local_method_-name	Specify the local method
			local_method_-pointer	Specify the local method
	Optional		local_search_-probability	Probability of executing local searches

Description

In the `embedded` approach, a tightly-coupled hybrid is employed in which a subordinate local method provides periodic refinements to a top-level global method.

Global and local method strings supplied with the `global_method_pointer` and `local_method_pointer` specifications identify the two methods to be used. The `local_search_probability` setting is an optional specification for supplying the probability (between 0.0 and 1.0) of employing local search to improve estimates within the global search.

global_method_name

- Keywords Area
- method
- hybrid
- embedded
- global_method_name

Specify the global method

Topics

This keyword is related to the topics:

- block_pointer
- not_yet_reviewed

Specification

Alias: none

Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		global_model_-pointer	

Description

See parent page.

global_model_pointer

- Keywords Area
- method
- hybrid
- embedded
- global_method_name

- global_model_pointer

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): STRING

global_method_pointer

- Keywords Area
- method
- hybrid
- embedded
- global_method_pointer

Specify the global method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

See parent page.

local_method_name

- Keywords Area
- method
- hybrid
- embedded
- local_method_name

Specify the local method

Topics

This keyword is related to the topics:

- block_pointer
- not_yet_reviewed

Specification

Alias: none

Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		local_model_-pointer	

Description

See parent page.

local_model_pointer

- Keywords Area
- method
- hybrid
- embedded
- local_method_name
- local_model_pointer

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none

Argument(s): STRING

local_method_pointer

- Keywords Area
- method
- hybrid
- embedded
- local_method_pointer

Specify the local method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

See parent page.

local_search_probability

- Keywords Area
- method
- hybrid
- embedded
- local_search_probability

Probability of executing local searches

Specification

Alias: none
Argument(s): REAL

Description

In the `embedded` approach, global and local method strings supplied with the `global_method_pointer` and `local_method_pointer` specifications identify the two methods to be used. As for other meta-iterators, method pointers can be replaced by method names and optional model pointers. The `local_search_probability` setting is an optional specification for supplying the probability (between 0.0 and 1.0) of employing local search to improve estimates within the global search.

collaborative

- Keywords Area
- method
- hybrid
- collaborative

Multiple methods run concurrently and share information

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required(*Choose One*)	**Group 1**	method_name_list	
			method_pointer_list	

Description

In the collaborative approach, multiple methods work together and share solutions while executing concurrently.

method_name_list

- Keywords Area
- method
- hybrid
- collaborative
- method_name_list

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- block_pointer
- not_yet_reviewed

Specification

Alias: none

Argument(s): STRINGLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		model_pointer_list	

model_pointer_list

- Keywords Area
- method
- hybrid
- collaborative
- method_name_list
- model_pointer_list

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- block_pointer
- not_yet_reviewed

Specification

Alias: none
Argument(s): STRING

method_pointer_list

- Keywords Area
- method
- hybrid
- collaborative
- method_pointer_list

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- block_pointer
- not_yet_reviewed

Specification

Alias: none
Argument(s): STRINGLIST

iterator_servers

- Keywords Area
- method
- hybrid
- iterator_servers

Specify the number of iterator servers when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `iterator_servers` specification supports user override of the automatic parallel configuration for the number of iterator servers. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired number of partitions at the iterator parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

iterator_scheduling

- Keywords Area
- method
- hybrid
- iterator_scheduling

Specify the scheduling of concurrent iterators when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism
- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	master	Specify a dedicated master partition for parallel iterator scheduling
			peer	Specify a peer partition for parallel iterator scheduling

Description

When Dakota is run in parallel, the partition type of the iterator servers is determined automatically. If this setting is undesirable, it may be overridden by the user using the `iterator_scheduling` keyword.

master

- Keywords Area
- method
- hybrid
- iterator_scheduling
- master

Specify a dedicated master partition for parallel iterator scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none

Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a dedicated master partition. In a dedicated master partition, one processor (the ”master”) dynamically schedules work on the iterator servers. This reduces the number of processors available to create servers by 1.

peer

- Keywords Area
- method
- hybrid
- iterator_scheduling
- peer

Specify a peer partition for parallel iterator scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a peer partition. In a peer partition, all processors are available to be assigned to iterator servers. Note that unlike the case of `evaluation_scheduling`, it is not possible to specify `static` or `dynamic`.

processors_per_iterator

- Keywords Area
- method
- hybrid
- processors_per_iterator

Specify the number of processors per iterator server when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `processors_per_iterator` specification supports user override of the automatic parallel configuration for the number of processors in each iterator server. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired server size at the iterator parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

5.2.11 multi_start

- Keywords Area
- method
- multi_start

Multi-Start Optimization Method

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	method_name	Name of a sub-method to apply to surrogate model
			method_pointer	Pointer to sub-method to apply to surrogate
	Optional		random_starts	Number of random starting points
	Optional		starting_points	List of user-specified starting points
	Optional		iterator_servers	Specify the number of iterator servers when Dakota is run in parallel
	Optional		iterator_scheduling	Specify the scheduling of concurrent iterators when Dakota is run in parallel

	Optional	processors_per_-iterator	Specify the number of processors per iterator server when Dakota is run in parallel

Description

In the multi-start iteration method (`multi_start`), a series of iterator runs are performed for different values of parameters in the model. A common use is for multi-start optimization (i.e., different local optimization runs from different starting points for the design variables), but the concept and the code are more general. Multi-start iteration is implemented within the MetaIterator branch of the Iterator hierarchy within the ConcurrentMeta-Iterator class. Additional information on the multi-start algorithm is available in the Users Manual [[4] "Adams et al., 2010"].

The `multi_start` meta-iterator must specify a sub-iterator using either a `method_pointer` or a `method_name` plus optional `model_pointer`. This iterator is responsible for completing a series of iterative analyses from a set of different starting points. These starting points can be specified as follows: (1) using `random_starts`, for which the specified number of starting points are selected randomly within the variable bounds, (2) using `starting_points`, in which the starting values are provided in a list, or (3) using both `random_starts` and `starting_points`, for which the combined set of points will be used. In aggregate, at least one starting point must be specified. The most common example of a multi-start algorithm is multi-start optimization, in which a series of optimizations are performed from different starting values for the design variables. This can be an effective approach for problems with multiple minima.

method_name

- Keywords Area
- method
- multi_start
- method_name

Name of a sub-method to apply to surrogate model

Topics

This keyword is related to the topics:

- block_pointer
- problem

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		model_pointer	Specify which model block will be used by a method

Description

The `surrogate_based_local` and `surrogate_based_global` methods must specify an optimization or least squares sub-method to apply to the surrogate model.

Option 2 is using `approx_method_name` (e.g., 'npsol_sqp') which identifies a sub-method for the sub-problem minimizer (which will use default minimizer settings).

model_pointer

- Keywords Area
- method
- multi_start
- method_name
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

method_pointer

- Keywords Area
- method
- multi_start
- method_pointer

Pointer to sub-method to apply to surrogate

Topics

This keyword is related to the topics:

- block_pointer
- problem

Specification

Alias: none
Argument(s): STRING

Description

The `surrogate_based_local` and `surrogate_based_global` methods must specify an optimization or least squares sub-method to apply to the surrogate model.

Option 1 is using `approx_method_pointer` (e.g., 'NLP1') which identifies a full sub-method specification for the sub-problem minimizer (which allows non-default minimizer settings).

Any `model_pointer` identified in an approximate sub-method specification is ignored. The surrogate model is pointed from the `model_pointer` keyword in the `surrogate_based_local` and `surrogate_based_global` `method` block.

random_starts

- Keywords Area
- method
- multi_start
- random_starts

Number of random starting points

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		seed	Specify the seed of the random number generator

Description

The `multi_start` meta-iterator must specify a sub-iterator using either a `method_pointer` or a `method_name` plus optional `model_pointer`. This iterator is responsible for completing a series of iterative analyses from a set of different starting points. These starting points can be specified as follows: (1) using `random_starts`, for which the specified number of starting points are selected randomly within the variable bounds, (2) using `starting_points`, in which the starting values are provided in a list, or (3) using both `random_starts` and `starting_points`, for which the combined set of points will be used.

seed

- Keywords Area
- method
- multi_start
- random_starts
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

starting_points

- Keywords Area
- method
- multi_start
- starting_points

List of user-specified starting points

Specification

Alias: none
Argument(s): REALLIST

Description

The `multi_start` meta-iterator must specify a sub-iterator using either a `method_pointer` or a `method_name` plus optional `model_pointer`. This iterator is responsible for completing a series of iterative analyses from a set of different starting points. These starting points can be specified as follows: (1) using `random_starts`, for which the specified number of starting points are selected randomly within the variable bounds, (2) using `starting_points`, in which the starting values are provided in a list, or (3) using both `random_starts` and `starting_points`, for which the combined set of points will be used.

iterator_servers

- Keywords Area
- method
- multi_start
- iterator_servers

Specify the number of iterator servers when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `iterator_servers` specification supports user override of the automatic parallel configuration for the number of iterator servers. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired number of partitions at the iterator parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

iterator_scheduling

- Keywords Area
- method
- multi_start
- iterator_scheduling

Specify the scheduling of concurrent iterators when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	master	Specify a dedicated master partition for parallel iterator scheduling
			peer	Specify a peer partition for parallel iterator scheduling

Description

When Dakota is run in parallel, the partition type of the iterator servers is determined automatically. If this setting is undesirable, it may be overridden by the user using the `iterator_scheduling` keyword.

master

- Keywords Area
- method
- multi_start
- iterator_scheduling
- master

Specify a dedicated master partition for parallel iterator scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a dedicated master partition. In a dedicated master partition, one processor (the "master") dynamically schedules work on the iterator servers. This reduces the number of processors available to create servers by 1.

peer

- Keywords Area
- method
- multi_start
- iterator_scheduling
- peer

Specify a peer partition for parallel iterator scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a peer partition. In a peer partition, all processors are available to be assigned to iterator servers. Note that unlike the case of `evaluation_scheduling`, it is not possible to specify `static` or `dynamic`.

processors_per_iterator

- Keywords Area
- method
- multi_start
- processors_per_iterator

Specify the number of processors per iterator server when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `processors_per_iterator` specification supports user override of the automatic parallel configuration for the number of processors in each iterator server. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired server size at the iterator parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

5.2.12 pareto_set

- Keywords Area
- method
- pareto_set

Pareto set optimization

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	method_name	Name of a sub-method to apply to surrogate model
			method_pointer	Pointer to sub-method to apply to surrogate
	Optional		random_weight_-sets	Number of random weighting sets
	Optional		weight_sets	List of user-specified weighting sets
	Optional		iterator_servers	Specify the number of iterator servers when Dakota is run in parallel
	Optional		iterator_scheduling	Specify the scheduling of concurrent iterators when Dakota is run in parallel

	Optional	processors_per_-iterator	Specify the number of processors per iterator server when Dakota is run in parallel

Description

In the pareto set minimization method (`pareto_set`), a series of optimization or least squares calibration runs are performed for different weightings applied to multiple objective functions. This set of optimal solutions defines a "Pareto set," which is useful for investigating design trade-offs between competing objectives. The code is similar enough to the `multi_start` technique that both algorithms are implemented in the same Concurrent-MetaIterator class.

The `pareto_set` specification must identify an optimization or least squares calibration method using either a `method_pointer` or a `method_name` plus optional `model_pointer`. This minimizer is responsible for computing a set of optimal solutions from a set of response weightings (multi-objective weights or least squares term weights). These weightings can be specified as follows: (1) using `random_weight_sets`, in which case weightings are selected randomly within [0,1] bounds, (2) using `weight_sets`, in which the weighting sets are specified in a list, or (3) using both `random_weight_sets` and `weight_sets`, for which the combined set of weights will be used. In aggregate, at least one set of weights must be specified. The set of optimal solutions is called the "pareto set," which can provide valuable design trade-off information when there are competing objectives.

method_name

- Keywords Area
- method
- pareto_set
- method_name

Name of a sub-method to apply to surrogate model

Topics

This keyword is related to the topics:

- block_pointer
- problem

Specification

Alias: opt_method_name
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		model_pointer	Specify which model block will be used by a method

Description

The `surrogate_based_local` and `surrogate_based_global` methods must specify an optimization or least squares sub-method to apply to the surrogate model.

Option 2 is using `approx_method_name` (e.g., 'npsol_sqp') which identifies a sub-method for the sub-problem minimizer (which will use default minimizer settings).

model_pointer

- Keywords Area
- method
- pareto_set
- method_name
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: opt_model_pointer
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

method_pointer

- Keywords Area
- method
- pareto_set
- method_pointer

Pointer to sub-method to apply to surrogate

Topics

This keyword is related to the topics:

- block_pointer
- problem

Specification

Alias: opt_method_pointer
Argument(s): STRING

Description

The `surrogate_based_local` and `surrogate_based_global` methods must specify an optimization or least squares sub-method to apply to the surrogate model.

Option 1 is using `approx_method_pointer` (e.g., 'NLP1') which identifies a full sub-method specification for the sub-problem minimizer (which allows non-default minimizer settings).

Any `model_pointer` identified in an approximate sub-method specification is ignored. The surrogate model is pointed from the `model_pointer` keyword in the `surrogate_based_local` and `surrogate_based_global` `method` block.

random_weight_sets

- Keywords Area
- method
- pareto_set
- random_weight_sets

Number of random weighting sets

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		seed	Specify the seed of the random number generator

Description

The `pareto_set` specification must identify an optimization or least squares calibration method using either a `method_pointer` or a `method_name` plus optional `model_pointer`. This minimizer is responsible for computing a set of optimal solutions from a set of response weightings (multi-objective weights or least squares term weights). These weightings can be specified as follows: (1) using `random_weight_sets`, in which case weightings are selected randomly within [0,1] bounds, (2) using `weight_sets`, in which the weighting sets are

specified in a list, or (3) using both `random_weight_sets` and `weight_sets`, for which the combined set of weights will be used. In aggregate, at least one set of weights must be specified. The set of optimal solutions is called the "pareto set," which can provide valuable design trade-off information when there are competing objectives.

seed

- Keywords Area
- method
- pareto_set
- random_weight_sets
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

weight_sets

- Keywords Area
- method
- pareto_set
- weight_sets

List of user-specified weighting sets

Specification

Alias: multi_objective_weight_sets
Argument(s): REALLIST

Description

The `pareto_set` specification must identify an optimization or least squares calibration method using either a `method_pointer` or a `method_name` plus optional `model_pointer`. This minimizer is responsible for computing a set of optimal solutions from a set of response weightings (multi-objective weights or least squares term weights). These weightings can be specified as follows: (1) using `random_weight_sets`, in which case weightings are selected randomly within [0,1] bounds, (2) using `weight_sets`, in which the weighting sets are specified in a list, or (3) using both `random_weight_sets` and `weight_sets`, for which the combined set of weights will be used. In aggregate, at least one set of weights must be specified. The set of optimal solutions is called the "pareto set," which can provide valuable design trade-off information when there are competing objectives.

iterator_servers

- Keywords Area
- method
- pareto_set
- iterator_servers

Specify the number of iterator servers when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `iterator_servers` specification supports user override of the automatic parallel configuration for the number of iterator servers. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired number of partitions at the iterator parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

iterator_scheduling

- Keywords Area
- method
- pareto_set
- iterator_scheduling

Specify the scheduling of concurrent iterators when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	master	Specify a dedicated master partition for parallel iterator scheduling
			peer	Specify a peer partition for parallel iterator scheduling

Description

When Dakota is run in parallel, the partition type of the iterator servers is determined automatically. If this setting is undesirable, it may be overridden by the user using the `iterator_scheduling` keyword.

master

- Keywords Area
- method
- pareto_set
- iterator_scheduling
- master

Specify a dedicated master partition for parallel iterator scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a dedicated master partition. In a dedicated master partition, one processor (the ”master”) dynamically schedules work on the iterator servers. This reduces the number of processors available to create servers by 1.

peer

- Keywords Area
- method
- pareto_set
- iterator_scheduling
- peer

Specify a peer partition for parallel iterator scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a peer partition. In a peer partition, all processors are available to be assigned to iterator servers. Note that unlike the case of `evaluation_scheduling`, it is not possible to specify `static` or `dynamic`.

processors_per_iterator

- Keywords Area
- method
- pareto_set
- processors_per_iterator

Specify the number of processors per iterator server when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `processors_per_iterator` specification supports user override of the automatic parallel configuration for the number of processors in each iterator server. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired server size at the iterator parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

5.2.13 surrogate_based_local

- Keywords Area
- method
- surrogate_based_local

Local Surrogate Based Optimization

Topics

This keyword is related to the topics:

- surrogate_based_optimization_methods

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	method_pointer	Pointer to sub-method to apply to surrogate
			method_name	Name of a sub-method to apply to surrogate model
	Required		model_pointer	Specify which model block will be used by a method

	Optional	soft_convergence_-limit	Limit number of iterations w/ little improvement
	Optional	truth_surrogate_-bypass	Bypass lower level surrogates when performing truth verifications on a top level surrogate
	Optional	trust_region	Use trust region search method
	Optional	approx_-subproblem	Identify functions to be included in surrogate merit function
	Optional	merit_function	Select type of penalty or merit function
	Optional	acceptance_logic	Set criteria for trusted surrogate
	Optional	constraint_relax	Enable constraint relaxation

Description

In addition to the method independent controls for `max_iterations` and `convergence_tolerance`

For surrogate_based_local problems with nonlinear constraints, a number of algorithm formulations exist as described in [[22] "Eldred and Dunlavy, 2006"] and as summarized in the Advanced Examples section of the Models chapter of the Users Manual [[4] "Adams et al., 2010"].

Theory

In surrogate-based optimization (SBO) and surrogate-based nonlinear least squares (SBNLS), minimization occurs using a set of one or more approximations, defined from a surrogate model, that are built and periodically updated using data from a "truth" model. The surrogate model can be a global data fit (e.g., regression or interpolation of data generated from a design of computer experiments), a multipoint approximation, a local Taylor Series expansion, or a model hierarchy approximation (e.g., a low-fidelity simulation model), whereas the truth model involves a high-fidelity simulation model. The goals of surrogate-based methods are to reduce the total number of truth model simulations and, in the case of global data fit surrogates, to smooth noisy data with an easily navigated analytic function.

In the surrogate-based local method, a trust region approach is used to manage the minimization process to maintain acceptable accuracy between the surrogate model and the truth model (by limiting the range over which the surrogate model is trusted). The process involves a sequence of minimizations performed on the surrogate model and bounded by the trust region. At the end of each approximate minimization, the candidate optimum point is validated using the truth model. If sufficient decrease has been obtained in the truth model, the trust region is re-centered around the candidate optimum point and the trust region will either shrink, expand, or remain the same size depending on the accuracy with which the surrogate model predicted the truth model decrease. If sufficient decrease has not been attained, the trust region center is not updated and the entire trust region shrinks by a user-specified factor. The cycle then repeats with the construction of a new surrogate model, a minimization, and another test for sufficient decrease in the truth model. This cycle continues until convergence is attained.

See Also

These keywords may also be of interest:

- efficient_global
- surrogate_based_global

method_pointer

- Keywords Area
- method
- surrogate_based_local
- method_pointer

Pointer to sub-method to apply to surrogate

Topics

This keyword is related to the topics:

- block_pointer
- problem

Specification

Alias: approx_method_pointer
Argument(s): STRING

Description

The `surrogate_based_local` and `surrogate_based_global` methods must specify an optimization or least squares sub-method to apply to the surrogate model.

Option 1 is using `approx_method_pointer` (e.g., 'NLP1') which identifies a full sub-method specification for the sub-problem minimizer (which allows non-default minimizer settings).

Any `model_pointer` identified in an approximate sub-method specification is ignored. The surrogate model is pointed from the `model_pointer` keyword in the `surrogate_based_local` and `surrogate_based_global` `method` block.

method_name

- Keywords Area
- method
- surrogate_based_local
- method_name

Name of a sub-method to apply to surrogate model

Topics

This keyword is related to the topics:

- block_pointer
- problem

Specification

Alias: approx_method_name
Argument(s): STRING

Description

The `surrogate_based_local` and `surrogate_based_global` methods must specify an optimization or least squares sub-method to apply to the surrogate model.

Option 2 is using `approx_method_name` (e.g., 'npsol_sqp') which identifies a sub-method for the sub-problem minimizer (which will use default minimizer settings).

model_pointer

- Keywords Area
- method
- surrogate_based_local
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: approx_model_pointer
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

soft_convergence_limit

- Keywords Area
- method
- surrogate_based_local
- soft_convergence_limit

Limit number of iterations w/ little improvement

Specification

Alias: none
Argument(s): INTEGER

Description

`soft_convergence_limit` (a soft convergence control for the surrogate_based_local iterations which limits the number of consecutive iterations with improvement less than the convergence tolerance)

truth_surrogate_bypass

- Keywords Area
- method
- surrogate_based_local
- truth_surrogate_bypass

Bypass lower level surrogates when performing truth verifications on a top level surrogate

Specification

Alias: none
Argument(s): none

Description

`truth_surrogate_bypass` (a flag for bypassing all lower level surrogates when performing truth verifications on a top level surrogate).

trust_region

- Keywords Area
- method
- surrogate_based_local
- trust_region

Use trust region search method

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_size	Trust region initial size (relative to bounds)
	Optional		minimum_size	Trust region minimum size
	Optional		contract_threshold	Shrink trust region if trust region ratio is below this value
	Optional		expand_threshold	Expand trust region if trust region ratio is above this value
	Optional		contraction_factor	Amount by which step length is rescaled
	Optional		expansion_factor	Trust region expansion factor

Description

The `trust_region` optional group specification can be used to specify the initial size of the trust region (using `initial_size`) relative to the total variable bounds, the minimum size of the trust region (using `minimum_size`), the contraction factor for the trust region size (using `contraction_factor`) used when the surrogate model is performing poorly, and the expansion factor for the trust region size (using `expansion_factor`) used when the the surrogate model is performing well. Two additional commands are the trust region size contraction threshold (using `contract_threshold`) and the trust region size expansion threshold (using `expand_threshold`). These two commands are related to what is called the trust region ratio, which is the actual decrease in the truth model divided by the predicted decrease in the truth model in the current trust region. The command `contract_threshold` sets the minimum acceptable value for the trust region ratio, i.e., values below this threshold cause the trust region to shrink for the next surrogate_based_local iteration. The command `expand_threshold` determines the trust region value above which the trust region will expand for the next surrogate_based_local iteration.

initial_size

- Keywords Area
- method
- surrogate_based_local
- trust_region
- initial_size

Trust region initial size (relative to bounds)

Specification

Alias: none
Argument(s): REAL

Description

The trust_region optional group specification can be used to specify the initial size of the trust region (using initial_size) relative to the total variable bounds, the minimum size of the trust region (using minimum_size), the contraction factor for the trust region size (using contraction_factor) used when the surrogate model is performing poorly, and the expansion factor for the trust region size (using expansion_factor) used when the the surrogate model is performing well. Two additional commands are the trust region size contraction threshold (using contract_threshold) and the trust region size expansion threshold (using expand_threshold). These two commands are related to what is called the trust region ratio, which is the actual decrease in the truth model divided by the predicted decrease in the truth model in the current trust region. The command contract_threshold sets the minimum acceptable value for the trust region ratio, i.e., values below this threshold cause the trust region to shrink for the next SBL iteration. The command expand_threshold determines the trust region value above which the trust region will expand for the next SBL iteration.

minimum_size

- Keywords Area
- method
- surrogate_based_local
- trust_region
- minimum_size

Trust region minimum size

Specification

Alias: none
Argument(s): REAL

Description

The trust_region optional group specification can be used to specify the initial size of the trust region (using initial_size) relative to the total variable bounds, the minimum size of the trust region (using minimum_size), the contraction factor for the trust region size (using contraction_factor) used when the surrogate model is performing poorly, and the expansion factor for the trust region size (using expansion_factor) used when the the surrogate model is performing well. Two additional commands are the trust region size contraction threshold (using contract_threshold) and the trust region size expansion threshold (using expand_threshold). These two commands are related to what is called the trust region ratio, which is the actual decrease in the truth model divided by the predicted decrease in the truth model in the current trust region. The command contract_threshold sets the minimum acceptable value for the trust region ratio, i.e., values below this threshold cause the trust region to shrink for the next SBL iteration. The command expand_threshold determines the trust region value above which the trust region will expand for the next SBL iteration.

contract_threshold

- Keywords Area
- method
- surrogate_based_local
- trust_region
- contract_threshold

Shrink trust region if trust region ratio is below this value

Specification

Alias: none
Argument(s): REAL

Description

The `trust_region` optional group specification can be used to specify the initial size of the trust region (using `initial_size`) relative to the total variable bounds, the minimum size of the trust region (using `minimum_size`), the contraction factor for the trust region size (using `contraction_factor`) used when the surrogate model is performing poorly, and the expansion factor for the trust region size (using `expansion_factor`) used when the the surrogate model is performing well. Two additional commands are the trust region size contraction threshold (using `contract_threshold`) and the trust region size expansion threshold (using `expand_threshold`). These two commands are related to what is called the trust region ratio, which is the actual decrease in the truth model divided by the predicted decrease in the truth model in the current trust region. The command `contract_threshold` sets the minimum acceptable value for the trust region ratio, i.e., values below this threshold cause the trust region to shrink for the next SBL iteration. The command `expand_threshold` determines the trust region value above which the trust region will expand for the next SBL iteration.

expand_threshold

- Keywords Area
- method
- surrogate_based_local
- trust_region
- expand_threshold

Expand trust region if trust region ratio is above this value

Specification

Alias: none
Argument(s): REAL

Description

The trust_region optional group specification can be used to specify the initial size of the trust region (using initial_size) relative to the total variable bounds, the minimum size of the trust region (using minimum_size), the contraction factor for the trust region size (using contraction_factor) used when the surrogate model is performing poorly, and the expansion factor for the trust region size (using expansion_factor) used when the the surrogate model is performing well. Two additional commands are the trust region size contraction threshold (using contract_threshold) and the trust region size expansion threshold (using expand_threshold). These two commands are related to what is called the trust region ratio, which is the actual decrease in the truth model divided by the predicted decrease in the truth model in the current trust region. The command contract_threshold sets the minimum acceptable value for the trust region ratio, i.e., values below this threshold cause the trust region to shrink for the next SBL iteration. The command expand_threshold determines the trust region value above which the trust region will expand for the next SBL iteration.

contraction_factor

- Keywords Area
- method
- surrogate_based_local
- trust_region
- contraction_factor

Amount by which step length is rescaled

Specification

Alias: none
Argument(s): REAL

Description

For pattern search methods, contraction_factor specifies the amount by which step length is rescaled after unsuccesful iterates, must be strictly between 0 and 1.

For methods that can expand the step length, the expansion is 1/ contraction_factor

expansion_factor

- Keywords Area
- method
- surrogate_based_local
- trust_region
- expansion_factor

Trust region expansion factor

Specification

Alias: none
Argument(s): REAL

Description

The `trust_region` optional group specification can be used to specify the initial size of the trust region (using `initial_size`) relative to the total variable bounds, the minimum size of the trust region (using `minimum_size`), the contraction factor for the trust region size (using `contraction_factor`) used when the surrogate model is performing poorly, and the expansion factor for the trust region size (using `expansion_factor`) used when the the surrogate model is performing well. Two additional commands are the trust region size contraction threshold (using `contract_threshold`) and the trust region size expansion threshold (using `expand_threshold`). These two commands are related to what is called the trust region ratio, which is the actual decrease in the truth model divided by the predicted decrease in the truth model in the current trust region. The command `contract_threshold` sets the minimum acceptable value for the trust region ratio, i.e., values below this threshold cause the trust region to shrink for the next SBL iteration. The command `expand_threshold` determines the trust region value above which the trust region will expand for the next SBL iteration.

approx_subproblem

- Keywords Area
- method
- surrogate_based_local
- approx_subproblem

Identify functions to be included in surrogate merit function

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	original_primary	Construct approximations of all primary functions
			single_objective	Construct approximation a single objective functions only

			augmented_-lagrangian_-objective	Augmented Lagrangian approximate subproblem formulation
			lagrangian_-objective	Lagrangian approximate subproblem formulation
	Required*(Choose One)*	**Group 2**	original_-constraints	Use the constraints directly
			linearized_-constraints	Use linearized approximations to the constraints
			no_constraints	Don't use constraints

Description

First, the "primary" functions (that is, the objective functions or calibration terms) in the approximate subproblem can be selected to be surrogates of the original primary functions (`original_primary`), a single objective function (`single_objective`) formed from the primary function surrogates, or either an augmented Lagrangian merit function (`augmented_lagrangian_objective`) or a Lagrangian merit function (`lagrangian_-objective`) formed from the primary and secondary function surrogates. The former option may imply the use of a nonlinear least squares method, a multiobjective optimization method, or a single objective optimization method to solve the approximate subproblem, depending on the definition of the primary functions. The latter three options all imply the use of a single objective optimization method regardless of primary function definition. Second, the surrogate constraints in the approximate subproblem can be selected to be surrogates of the original constraints (`original_constraints`) or linearized approximations to the surrogate constraints (`linearized_constraints`), or constraints can be omitted from the subproblem (`no_constraints`).

original_primary

- Keywords Area
- method
- surrogate_based_local
- approx_subproblem
- original_primary

Construct approximations of all primary functions

Specification

Alias: none

Argument(s): none

Description

For SBL problems with nonlinear constraints, a number of algorithm formulations exist as described in [[22] "Eldred and Dunlavy, 2006"] and as summarized in the Advanced Examples section of the Models chapter of the Users Manual [[4] "Adams et al., 2010"]. First, the "primary" functions (that is, the objective functions or calibration terms) in the approximate subproblem can be selected to be surrogates of the original primary functions (`original_primary`), a single objective function (`single_objective`) formed from the primary function surrogates, or either an augmented Lagrangian merit function (`augmented_lagrangian_objective`) or a Lagrangian merit function (`lagrangian_objective`) formed from the primary and secondary function surrogates. The former option may imply the use of a nonlinear least squares method, a multiobjective optimization method, or a single objective optimization method to solve the approximate subproblem, depending on the definition of the primary functions. The latter three options all imply the use of a single objective optimization method regardless of primary function definition.

single_objective

- Keywords Area
- method
- surrogate_based_local
- approx_subproblem
- single_objective

Construct approximation a single objective functions only

Specification

Alias: none
Argument(s): none

Description

For SBL problems with nonlinear constraints, a number of algorithm formulations exist as described in [[22] "Eldred and Dunlavy, 2006"] and as summarized in the Advanced Examples section of the Models chapter of the Users Manual [[4] "Adams et al., 2010"]. First, the "primary" functions (that is, the objective functions or calibration terms) in the approximate subproblem can be selected to be surrogates of the original primary functions (`original_primary`), a single objective function (`single_objective`) formed from the primary function surrogates, or either an augmented Lagrangian merit function (`augmented_lagrangian_objective`) or a Lagrangian merit function (`lagrangian_objective`) formed from the primary and secondary function surrogates. The former option may imply the use of a nonlinear least squares method, a multiobjective optimization method, or a single objective optimization method to solve the approximate subproblem, depending on the definition of the primary functions. The latter three options all imply the use of a single objective optimization method regardless of primary function definition.

augmented_lagrangian_objective

- Keywords Area
- method

- surrogate_based_local
- approx_subproblem
- augmented_lagrangian_objective

Augmented Lagrangian approximate subproblem formulation

Specification

Alias: none
Argument(s): none

Description

For SBL problems with nonlinear constraints, a number of algorithm formulations exist as described in [[22] "Eldred and Dunlavy, 2006"] and as summarized in the Advanced Examples section of the Models chapter of the Users Manual [[4] "Adams et al., 2010"]. First, the "primary" functions (that is, the objective functions or calibration terms) in the approximate subproblem can be selected to be surrogates of the original primary functions (`original_primary`), a single objective function (`single_objective`) formed from the primary function surrogates, or either an augmented Lagrangian merit function (`augmented_lagrangian_objective`) or a Lagrangian merit function (`lagrangian_objective`) formed from the primary and secondary function surrogates. The former option may imply the use of a nonlinear least squares method, a multiobjective optimization method, or a single objective optimization method to solve the approximate subproblem, depending on the definition of the primary functions. The latter three options all imply the use of a single objective optimization method regardless of primary function definition.

lagrangian_objective

- Keywords Area
- method
- surrogate_based_local
- approx_subproblem
- lagrangian_objective

Lagrangian approximate subproblem formulation

Specification

Alias: none
Argument(s): none

Description

For SBL problems with nonlinear constraints, a number of algorithm formulations exist as described in [[22] "Eldred and Dunlavy, 2006"] and as summarized in the Advanced Examples section of the Models chapter of the Users Manual [[4] "Adams et al., 2010"]. First, the "primary" functions (that is, the objective functions or calibration terms) in the approximate subproblem can be selected to be surrogates of the original primary functions (`original_primary`), a single objective function (`single_objective`) formed from the primary function

surrogates, or either an augmented Lagrangian merit function (`augmented_lagrangian_objective`) or a Lagrangian merit function (`lagrangian_objective`) formed from the primary and secondary function surrogates. The former option may imply the use of a nonlinear least squares method, a multiobjective optimization method, or a single objective optimization method to solve the approximate subproblem, depending on the definition of the primary functions. The latter three options all imply the use of a single objective optimization method regardless of primary function definition.

original_constraints

- Keywords Area
- method
- surrogate_based_local
- approx_subproblem
- original_constraints

Use the constraints directly

Specification

Alias: none
Argument(s): none

Description

The surrogate constraints in the approximate subproblem can be selected to be surrogates of the original constraints (`original_constraints`) or linearized approximations to the surrogate constraints (`linearized_constraints`), or constraints can be omitted from the subproblem (`no_constraints`).

linearized_constraints

- Keywords Area
- method
- surrogate_based_local
- approx_subproblem
- linearized_constraints

Use linearized approximations to the constraints

Specification

Alias: none
Argument(s): none

Description

The surrogate constraints in the approximate subproblem can be selected to be surrogates of the original constraints (`original_constraints`) or linearized approximations to the surrogate constraints (`linearized_constraints`), or constraints can be omitted from the subproblem (`no_constraints`).

no_constraints

- Keywords Area
- method
- surrogate_based_local
- approx_subproblem
- no_constraints

Don't use constraints

Specification

Alias: none
Argument(s): none

Description

The surrogate constraints in the approximate subproblem can be selected to be surrogates of the original constraints (`original_constraints`) or linearized approximations to the surrogate constraints (`linearized_constraints`), or constraints can be omitted from the subproblem (`no_constraints`).

merit_function

- Keywords Area
- method
- surrogate_based_local
- merit_function

Select type of penalty or merit function

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Required*(Choose One)*	**Group 1**	penalty_merit	Use penalty merit function
			adaptive_penalty_-merit	Use adaptive penalty merit function
			lagrangian_merit	Use first-order Lagrangian merit function
			augmented_-lagrangian_merit	Use combined penalty and zeroth-order Lagrangian merit function

Description

Following optimization of the approximate subproblem, the candidate iterate is evaluated using a merit function, which can be selected to be a simple penalty function with penalty ramped by surrogate_based_local iteration number (`penalty_merit`), an adaptive penalty function where the penalty ramping may be accelerated in order to avoid rejecting good iterates which decrease the constraint violation (`adaptive_penalty_merit`), a Lagrangian merit function which employs first-order Lagrange multiplier updates (`lagrangian_merit`), or an augmented Lagrangian merit function which employs both a penalty parameter and zeroth-order Lagrange multiplier updates (`augmented_lagrangian_merit`). When an augmented Lagrangian is selected for either the subproblem objective or the merit function (or both), updating of penalties and multipliers follows the approach described in [[15] "Conn et al., 2000"].

penalty_merit

- Keywords Area
- method
- surrogate_based_local
- merit_function
- penalty_merit

Use penalty merit function

Specification

Alias: none
Argument(s): none

Description

Second, the surrogate constraints in the approximate subproblem can be selected to be surrogates of the original constraints (`original_constraints`) or linearized approximations to the surrogate constraints (`linearized_constraints`), or constraints can be omitted from the subproblem (`no_constraints`). Following optimization of the approximate subproblem, the candidate iterate is evaluated using a merit function, which can be selected to be a simple penalty function with penalty ramped by SBL iteration number (`penalty_merit`), an

adaptive penalty function where the penalty ramping may be accelerated in order to avoid rejecting good iterates which decrease the constraint violation (adaptive_penalty_merit), a Lagrangian merit function which employs first-order Lagrange multiplier updates (lagrangian_merit), or an augmented Lagrangian merit function which employs both a penalty parameter and zeroth-order Lagrange multiplier updates (augmented_lagrangian_merit). When an augmented Lagrangian is selected for either the subproblem objective or the merit function (or both), updating of penalties and multipliers follows the approach described in [[15] "Conn et al., 2000"].

adaptive_penalty_merit

- Keywords Area
- method
- surrogate_based_local
- merit_function
- adaptive_penalty_merit

Use adaptive penalty merit function

Specification

Alias: none
Argument(s): none

Description

Second, the surrogate constraints in the approximate subproblem can be selected to be surrogates of the original constraints (original_constraints) or linearized approximations to the surrogate constraints (linearized_constraints), or constraints can be omitted from the subproblem (no_constraints). Following optimization of the approximate subproblem, the candidate iterate is evaluated using a merit function, which can be selected to be a simple penalty function with penalty ramped by SBL iteration number (penalty_merit), an adaptive penalty function where the penalty ramping may be accelerated in order to avoid rejecting good iterates which decrease the constraint violation (adaptive_penalty_merit), a Lagrangian merit function which employs first-order Lagrange multiplier updates (lagrangian_merit), or an augmented Lagrangian merit function which employs both a penalty parameter and zeroth-order Lagrange multiplier updates (augmented_lagrangian_merit). When an augmented Lagrangian is selected for either the subproblem objective or the merit function (or both), updating of penalties and multipliers follows the approach described in [[15] "Conn et al., 2000"].

lagrangian_merit

- Keywords Area
- method
- surrogate_based_local
- merit_function
- lagrangian_merit

Use first-order Lagrangian merit function

Specification

Alias: none
Argument(s): none

Description

Second, the surrogate constraints in the approximate subproblem can be selected to be surrogates of the original constraints (`original_constraints`) or linearized approximations to the surrogate constraints (`linearized_constraints`), or constraints can be omitted from the subproblem (`no_constraints`). Following optimization of the approximate subproblem, the candidate iterate is evaluated using a merit function, which can be selected to be a simple penalty function with penalty ramped by SBL iteration number (`penalty_merit`), an adaptive penalty function where the penalty ramping may be accelerated in order to avoid rejecting good iterates which decrease the constraint violation (`adaptive_penalty_merit`), a Lagrangian merit function which employs first-order Lagrange multiplier updates (`lagrangian_merit`), or an augmented Lagrangian merit function which employs both a penalty parameter and zeroth-order Lagrange multiplier updates (`augmented_lagrangian_merit`). When an augmented Lagrangian is selected for either the subproblem objective or the merit function (or both), updating of penalties and multipliers follows the approach described in [[15] "Conn et al., 2000"].

augmented_lagrangian_merit

- Keywords Area
- method
- surrogate_based_local
- merit_function
- augmented_lagrangian_merit

Use combined penalty and zeroth-order Lagrangian merit function

Specification

Alias: none
Argument(s): none

Description

Second, the surrogate constraints in the approximate subproblem can be selected to be surrogates of the original constraints (`original_constraints`) or linearized approximations to the surrogate constraints (`linearized_constraints`), or constraints can be omitted from the subproblem (`no_constraints`). Following optimization of the approximate subproblem, the candidate iterate is evaluated using a merit function, which can be selected to be a simple penalty function with penalty ramped by SBL iteration number (`penalty_merit`), an adaptive penalty function where the penalty ramping may be accelerated in order to avoid rejecting good iterates which decrease the constraint violation (`adaptive_penalty_merit`), a Lagrangian merit function which employs first-order Lagrange multiplier updates (`lagrangian_merit`), or an augmented Lagrangian merit function which employs both a penalty parameter and zeroth-order Lagrange multiplier updates (`augmented_lagrangian_merit`). When an augmented Lagrangian is selected for either the subproblem objective or the merit function (or both), updating of penalties and multipliers follows the approach described in [[15] "Conn et al., 2000"].

acceptance_logic

- Keywords Area
- method
- surrogate_based_local
- acceptance_logic

Set criteria for trusted surrogate

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** tr_ratio	**Dakota Keyword Description** Surrogate-Based Local iterate acceptance logic
			filter	Surrogate-Based Local iterate acceptance logic

Description

Following calculation of the merit function for the new iterate, the iterate is accepted or rejected and the trust region size is adjusted for the next surrogate_based_local iteration. Iterate acceptance is governed either by a trust region ratio (`tr_ratio`) formed from the merit function values or by a filter method (`filter`); however, trust region resizing logic is currently based only on the trust region ratio. For infeasible iterates, constraint relaxation can be used for balancing constraint satisfaction and progress made toward an optimum.

tr_ratio

- Keywords Area
- method
- surrogate_based_local
- acceptance_logic
- tr_ratio

Surrogate-Based Local iterate acceptance logic

Specification

Alias: none
Argument(s): none

Description

Following calculation of the merit function for the new iterate, the iterate is accepted or rejected and the trust region size is adjusted for the next SBL iteration. Iterate acceptance is governed either by a trust region ratio (`tr_ratio`) formed from the merit function values or by a filter method (`filter`); however, trust region resizing logic is currently based only on the trust region ratio. For infeasible iterates, constraint relaxation can be used for balancing constraint satisfaction and progress made toward an optimum. The command `constraint_relax` followed by a method name specifies the type of relaxation to be used. Currently, `homotopy` [[**?**] "Perez et al., 2004"] is the only available method for constraint relaxation, and this method is dependent on the presence of the NPSOL library within the Dakota executable.

filter

- Keywords Area
- method
- surrogate_based_local
- acceptance_logic
- filter

Surrogate-Based Local iterate acceptance logic

Specification

Alias: none
Argument(s): none

Description

Following calculation of the merit function for the new iterate, the iterate is accepted or rejected and the trust region size is adjusted for the next SBL iteration. Iterate acceptance is governed either by a trust region ratio (`tr_ratio`) formed from the merit function values or by a filter method (`filter`); however, trust region resizing logic is currently based only on the trust region ratio. For infeasible iterates, constraint relaxation can be used for balancing constraint satisfaction and progress made toward an optimum. The command `constraint_relax` followed by a method name specifies the type of relaxation to be used. Currently, `homotopy` [[**?**] "Perez et al., 2004"] is the only available method for constraint relaxation, and this method is dependent on the presence of the NPSOL library within the Dakota executable.

constraint_relax

- Keywords Area
- method
- surrogate_based_local

- constraint_relax

Enable constraint relaxation

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		homotopy	Surrogate-Based local constraint relaxation method for infeasible iterates

Description

The command `constraint_relax` followed by a method name specifies the type of relaxation to be used. Currently, `homotopy` [[69] "Perez et al., 2004"] is the only available method for constraint relaxation, and this method is dependent on the presence of the NPSOL library within the Dakota executable.

homotopy

- Keywords Area
- method
- surrogate_based_local
- constraint_relax
- homotopy

Surrogate-Based local constraint relaxation method for infeasible iterates

Specification

Alias: none
Argument(s): none

Description

Currently, `homotopy` [[**?**] "Perez et al., 2004"] is the only available method for constraint relaxation, and this method is dependent on the presence of the NPSOL library within the Dakota executable.

5.2.14 surrogate_based_global

- Keywords Area
- method
- surrogate_based_global

Global Surrogate Based Optimization

Topics

This keyword is related to the topics:

- surrogate_based_optimization_methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required(*Choose One*)	**Group 1**	method_pointer	Pointer to sub-method to apply to surrogate
			method_name	Name of a sub-method to apply to surrogate model
	Required		model_pointer	Specify which model block will be used by a method
	Optional		replace_points	(Recommended) Replace points in the surrogate training set, instead of appending

Description

The `surrogate_based_global` specification must identify:

- a sub-method, using either method_pointer or method_name
- model_pointer must be used to identify a surrogate model

`surrogate_based_global` works in an iterative scheme where optimization is performed on a global surrogate using the same bounds during each iteration.

- In one iteration, the optimal solutions of the surrogate model are found, and then a selected set of these optimal surrogate solutions are passed to the next iteration.
- At the next iteration, these surrogate points are evaluated with the "truth" model, and then these points are added back to the set of points upon which the next surrogate is constructed.

In this way, the optimization acts on a more accurate surrogate during each iteration, presumably driving to optimality quickly.

Method Independent Controls

- max_iterations is used as a stopping critierion (see note below)

Notes

We have some cautionary notes before using the surrogate-based global method:

- **This approach has no guarantee of convergence.**
- One might first try a single minimization method coupled with a surrogate model prior to using the surrogate-based global method. This is essentially equivalent to setting `max_iterations` to 1 and will allow one to get a sense of what surrogate types are the most accurate to use for the problem.
- Also note that one can specify that surrogates be built for all primary functions and constraints or for only a subset of these functions and constraints. This allows one to use a "truth" model directly for some of the response functions, perhaps due to them being much less expensive than other functions.
- We initially recommend a small number of maximum iterations, such as 3-5, to get a sense of how the optimization is evolving as the surrogate gets updated. If it appears to be changing significantly, then a larger number (used in combination with restart) may be needed.

Theory

In surrogate-based optimization (SBO) and surrogate-based nonlinear least squares (SBNLS), minimization occurs using a set of one or more approximations, defined from a surrogate model, that are built and periodically updated using data from a "truth" model. The surrogate model can be a global data fit (e.g., regression or interpolation of data generated from a design of computer experiments), a multipoint approximation, a local Taylor Series expansion, or a model hierarchy approximation (e.g., a low-fidelity simulation model), whereas the truth model involves a high-fidelity simulation model. The goals of surrogate-based methods are to reduce the total number of truth model simulations and, in the case of global data fit surrogates, to smooth noisy data with an easily navigated analytic function.

It was originally designed for MOGA (a multi-objective genetic algorithm). Since genetic algorithms often need thousands or tens of thousands of points to produce optimal or near-optimal solutions, the use of surrogates can be helpful for reducing the truth model evaluations. Instead of creating one set of surrogates for the individual objectives and running the optimization algorithm on the surrogate once, the idea is to select points along the (surrogate) Pareto frontier, which can be used to supplement the existing points.

In this way, one does not need to use many points initially to get a very accurate surrogate. The surrogate becomes more accurate as the iterations progress.

See Also

These keywords may also be of interest:

- efficient_global
- surrogate_based_local

method_pointer

- Keywords Area
- method
- surrogate_based_global
- method_pointer

Pointer to sub-method to apply to surrogate

Topics

This keyword is related to the topics:

- block_pointer
- problem

Specification

Alias: approx_method_pointer
Argument(s): STRING

Description

The `surrogate_based_local` and `surrogate_based_global` methods must specify an optimization or least squares sub-method to apply to the surrogate model.

Option 1 is using `approx_method_pointer` (e.g., 'NLP1') which identifies a full sub-method specification for the sub-problem minimizer (which allows non-default minimizer settings).

Any `model_pointer` identified in an approximate sub-method specification is ignored. The surrogate model is pointed from the `model_pointer` keyword in the `surrogate_based_local` and `surrogate_based_global` `method` block.

method_name

- Keywords Area
- method
- surrogate_based_global
- method_name

Name of a sub-method to apply to surrogate model

Topics

This keyword is related to the topics:

- block_pointer
- problem

Specification

Alias: approx_method_name
Argument(s): STRING

Description

The `surrogate_based_local` and `surrogate_based_global` methods must specify an optimization or least squares sub-method to apply to the surrogate model.

Option 2 is using `approx_method_name` (e.g., 'npsol_sqp') which identifies a sub-method for the sub-problem minimizer (which will use default minimizer settings).

model_pointer

- Keywords Area
- method
- surrogate_based_global
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: approx_model_pointer
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

replace_points

- Keywords Area
- method
- surrogate_based_global
- replace_points

(Recommended) Replace points in the surrogate training set, instead of appending

Specification

Alias: none
Argument(s): none

Description

The user has the option of appending the optimal points from the surrogate model to the current set of truth points or using the optimal points from the surrogate model to replace the optimal set of points from the previous iteration. Although appending to the set is the default behavior, at this time we strongly recommend using the option `replace_points` because it appears to be more accurate and robust.

5.2.15 dot_frcg

- Keywords Area
- method
- dot_frcg

A conjugate gradient optimization method

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional		linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional		linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional		linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional		linear_equality_-targets	Define target values for the linear equality constraints

	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

This is a duplicated keyword. Please use dot instead.

We here provide a caution regarding `dot_frcg`. In DOT Version 4.20, we have noticed inconsistent behavior of this algorithm across different versions of Linux. Our best assessment is that it is due to different treatments of uninitialized variables. As we do not know the intention of the code authors and maintaining DOT source code is outside of the Dakota project scope, we have not made nor are we recommending any code changes to address this. However, all users who use `dot_frcg` in DOT Version 4.20 should be aware that results may not be reliable.

See Also

These keywords may also be of interest:

- frcg

linear_inequality_constraint_matrix

- Keywords Area
- method
- dot_frcg
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- dot_frcg
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- dot_frcg
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- dot_frcg
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- dot_frcg
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- dot_frcg
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- dot_frcg
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- dot_frcg
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- dot_frcg
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$
$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- dot_frcg
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.16 dot_mmfd

- Keywords Area
- method
- dot_mmfd

Method of feasible directions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint

	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

This is a duplicated keyword. Please use dot instead.

See Also

These keywords may also be of interest:

- mmfd

linear_inequality_constraint_matrix

- Keywords Area
- method
- dot_mmfd
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- dot_mmfd
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `−bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `−DBL_MAX` < `−bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- dot_mmfd
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `−bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `−DBL_MAX` < `−bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- dot_mmfd
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- dot_mmfd
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- dot_mmfd
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- dot_mmfd
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- dot_mmfd
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- dot_mmfd
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- dot_mmfd
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.17 dot_bfgs

- Keywords Area
- method
- dot_bfgs

A conjugate gradient optimization method

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints

	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

This is a duplicated keyword. Please use dot instead.

See Also

These keywords may also be of interest:

- bfgs

linear_inequality_constraint_matrix

- Keywords Area
- method
- dot_bfgs

- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- dot_bfgs
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- dot_bfgs
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `−bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `−DBL_MAX` < `−bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- dot_bfgs
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- dot_bfgs
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- dot_bfgs
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- dot_bfgs
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- dot_bfgs
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified

- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- dot_bfgs
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$
$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- dot_bfgs
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.18 dot_slp

- Keywords Area
- method
- dot_slp

Sequential Linear Programming

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional		linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional		linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional		linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional		linear_equality_-targets	Define target values for the linear equality constraints
	Optional		linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional		linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional		model_pointer	Specify which model block will be used by a method

Description

This is a duplicated keyword. Please use dot instead.

See Also

These keywords may also be of interest:

- slp

linear_inequality_constraint_matrix

- Keywords Area
- method
- dot_slp
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- dot_slp
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- dot_slp
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- dot_slp
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- dot_slp
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- dot_slp
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- dot_slp
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- dot_slp
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- dot_slp
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- dot_slp
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.19 dot_sqp

- Keywords Area
- method
- dot_sqp

Sequential Quadratic Program

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional		linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional		linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional		linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional		linear_equality_-targets	Define target values for the linear equality constraints

	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

This is a duplicated keyword. Please use dot instead.

See Also

These keywords may also be of interest:

- sqp

linear_inequality_constraint_matrix

- Keywords Area
- method
- dot_sqp
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- dot_sqp
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- dot_sqp
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- dot_sqp
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- dot_sqp
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- dot_sqp
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- dot_sqp
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- dot_sqp
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- dot_sqp
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$
$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- dot_sqp
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.20 dot

- Keywords Area
- method
- dot

Access to methods in the DOT package

Topics

This keyword is related to the topics:

- package_dot

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	frcg	A conjugate gradient optimization method
			mmfd	Method of feasible directions
			bfgs	A conjugate gradient optimization method

			slp	Sequential Linear Programming
			sqp	Sequential Quadratic Program
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional		linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional		linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional		linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional		linear_equality_-targets	Define target values for the linear equality constraints
	Optional		linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional		linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional		model_pointer	Specify which model block will be used by a method

Description

The DOT library [[86] ”Vanderplaats Research and Development, 1995”] contains nonlinear programming optimizers, specifically the Broyden-Fletcher-Goldfarb-Shanno (Dakota’s `dot_bfgs` method) and Fletcher-Reeves conjugate gradient (Dakota’s `dot_frcg` method) methods for unconstrained optimization, and the modified method of feasible directions (Dakota’s `dot_mmfd` method), sequential linear programming (Dakota’s `dot_slp` method), and sequential quadratic programming (Dakota’s `dot_sqp` method) methods for constrained optimization.

Specialized handling of linear constraints is supported with DOT; linear constraint coefficients, bounds, and targets can be provided to DOT at start-up and tracked internally.

One of the five available methods in **Group 1** must be specified.

All these methods take the same **Optional Keywords** , dealing with linear equality and inequality constraints.

Method Independent Controls - Stopping Critiera

Stopping critiera are set by:

- max_iterations
- max_function_evaluations
- convergence_tolerance
- constraint_tolerance

Note: The `convergence_tolerance` criterion must be satisfied for two consecutive iterations before DOT will terminate.

Method Independent Controls - Output

The output verbosity specification controls the amount of information generated by DOT: the `silent` and `quiet` settings result in header information, final results, and objective function, constraint, and parameter information on each iteration; whereas the `verbose` and `debug` settings add additional information on gradients, search direction, one-dimensional search results, and parameter scaling factors.

Concurrency

DOT contains no parallel algorithms which can directly take advantage of concurrent evaluations. However, if `numerical_gradients` with `method_source dakota` is specified, then the finite difference function evaluations can be performed concurrently (using any of the parallel modes described in the Users Manual [[4] "Adams et al., 2010"]). In addition, if `speculative` is specified, then gradients (`dakota numerical` or `analytic` gradients) will be computed on each line search evaluation in order to balance the load and lower the total run time in parallel optimization studies.

frcg

- Keywords Area
- method
- dot
- frcg

A conjugate gradient optimization method

Topics

This keyword is related to the topics:

- package_dot

Specification

Alias: none

Argument(s): none

Description

We here provide a caution regarding `dot_frcg`. In DOT Version 4.20, we have noticed inconsistent behavior of this algorithm across different versions of Linux. Our best assessment is that it is due to different treatments of uninitialized variables. As we do not know the intention of the code authors and maintaining DOT source code is outside of the Dakota project scope, we have not made nor are we recommending any code changes to address this. However, all users who use `dot_frcg` in DOT Version 4.20 should be aware that results may not be reliable.

See package_dot for information related to all `DOT` methods.

See Also

These keywords may also be of interest:

- bfgs
- mmfd
- slp
- sqp

mmfd

- Keywords Area
- method
- dot
- mmfd

Method of feasible directions

Topics

This keyword is related to the topics:

- package_dot

Specification

Alias: none

Argument(s): none

Description

See package_dot for information related to all `DOT` methods.

See Also

These keywords may also be of interest:

- bfgs
- frcg
- slp
- sqp

bfgs

- Keywords Area
- method
- dot
- bfgs

A conjugate gradient optimization method

Topics

This keyword is related to the topics:

- package_dot

Specification

Alias: none
Argument(s): none

Description

See package_dot for information related to all `DOT` methods.

See Also

These keywords may also be of interest:

- frcg
- mmfd
- slp
- sqp

slp

- Keywords Area
- method
- dot
- slp

Sequential Linear Programming

Topics

This keyword is related to the topics:

- package_dot

Specification

Alias: none
Argument(s): none

Description

See package_dot for information related to all `DOT` methods.

See Also

These keywords may also be of interest:

- bfgs
- frcg
- mmfd
- sqp

sqp

- Keywords Area
- method
- dot
- sqp

Sequential Quadratic Program

Topics

This keyword is related to the topics:

- package_dot

Specification

Alias: none
Argument(s): none

Description

See package_dot for information related to all `DOT` methods.

See Also

These keywords may also be of interest:

- bfgs
- frcg
- mmfd
- slp

linear_inequality_constraint_matrix

- Keywords Area
- method
- dot
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- dot
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- dot
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- dot
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- dot
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- dot
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- dot
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of $0.$ for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- dot
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- dot
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- dot
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.21 conmin_frcg

- Keywords Area
- method
- conmin_frcg

A conjugate gradient optimization method

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional		linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional		linear_inequality_-scales	Define the characteristic values to scale linear inequalities

	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

This is a duplicated keyword. Please use conmin instead.

See Also

These keywords may also be of interest:

- frcg

linear_inequality_constraint_matrix

- Keywords Area
- method
- conmin_frcg
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- conmin_frcg
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- conmin_frcg
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- conmin_frcg
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- conmin_frcg
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- conmin_frcg
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- conmin_frcg
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- conmin_frcg
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- conmin_frcg
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$
$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- conmin_frcg
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.22 conmin_mfd

- Keywords Area
- method
- conmin_mfd

Method of feasible directions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint

	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

This is a duplicated keyword. Please use conmin instead.

See Also

These keywords may also be of interest:

- mfd

linear_inequality_constraint_matrix

- Keywords Area
- method
- conmin_mfd
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- conmin_mfd
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `−bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `−DBL_MAX` < `−bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- conmin_mfd
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `−bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `−DBL_MAX` < `−bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- conmin_mfd
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- conmin_mfd
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- conmin_mfd
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- conmin_mfd
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- conmin_mfd
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- conmin_mfd
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- conmin_mfd
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.23 conmin

- Keywords Area
- method
- conmin

Access to methods in the CONMIN library

Topics

This keyword is related to the topics:

- package_conmin

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	frcg	A conjugate gradient optimization method
			mfd	Method of feasible directions
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional		linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional		linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional		linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional		linear_equality_-targets	Define target values for the linear equality constraints
	Optional		linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional		linear_equality_-scales	Define the characteristic values to scale linear equalities

	Optional	model_pointer	Specify which model block will be used by a method

Description

The CONMIN library [[84] "Vanderplaats, 1973"] is a public domain library of nonlinear programming optimizers, specifically the Fletcher-Reeves conjugate gradient (Dakota's `conmin_frcg` method) method for unconstrained optimization, and the method of feasible directions (Dakota's `conmin_mfd` method) for constrained optimization. As CONMIN was a predecessor to the DOT commercial library, the algorithm controls are very similar.

One of the two available methods in **Group 1** must be specified.

All these methods take the same **Optional Keywords** , dealing with linear equality and inequality constraints.

See Also

These keywords may also be of interest:

- dot

frcg

- Keywords Area
- method
- conmin
- frcg

A conjugate gradient optimization method

Topics

This keyword is related to the topics:

- package_conmin

Specification

Alias: none

Argument(s): none

Description

The interpretations of the method independent controls for CONMIN are essentially identical to those for DOT.

See package_dot for information related to `CONMIN` methods.

See Also

These keywords may also be of interest:

- mfd
- frcg

mfd

- Keywords Area
- method
- conmin
- mfd

Method of feasible directions

Topics

This keyword is related to the topics:

- package_conmin

Specification

Alias: none
Argument(s): none

Description

The interpretations of the method independent controls for CONMIN are essentially identical to those for DOT. See package_dot for information related to `CONMIN` methods.

See Also

These keywords may also be of interest:

- frcg
- mmfd

linear_inequality_constraint_matrix

- Keywords Area
- method
- conmin
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- conmin
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- conmin
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- conmin
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- conmin
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- conmin
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- conmin
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- conmin
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified

- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- conmin
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- conmin
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.24 dl_solver

- Keywords Area
- method
- dl_solver

There is ZERO mention of this in the reference manual??

Topics

This keyword is related to the topics:

- problem
- not_yet_reviewed

Specification

Alias: none

Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional		linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional		linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional		linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional		linear_equality_-targets	Define target values for the linear equality constraints
	Optional		linear_equality_-scale_types	Specify how each linear equality constraint is scaled

	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

linear_inequality_constraint_matrix

- Keywords Area
- method
- dl_solver
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_-inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- dl_solver
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- dl_solver
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- dl_solver
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- dl_solver
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- dl_solver
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- dl_solver
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- dl_solver
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- dl_solver
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- dl_solver
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.25 npsol_sqp

- Keywords Area
- method
- npsol_sqp

Sequential Quadratic Program

Topics

This keyword is related to the topics:

- package_npsol
- sequential_quadratic_programming
- local_optimization_methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		verify_level	Verify the quality of analytic gradients
	Optional		function_precision	Specify the maximum precision of the analysis code responses
	Optional		linesearch_-tolerance	Choose how accurately the algorithm will compute the minimum in a line search
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints

	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

NPSOL provides an implementation of sequential quadratic programming that can be accessed with `npsol_sqp`.

Stopping Criteria

The method independent controls for `max_iterations` and `max_function_evaluations` limit the number of major SQP iterations and the number of function evaluations that can be performed during an NPSOL optimization. The `convergence_tolerance` control defines NPSOL's internal optimality tolerance which is used in evaluating if an iterate satisfies the first-order Kuhn-Tucker conditions for a minimum. The magnitude of `convergence_tolerance` approximately specifies the number of significant digits of accuracy desired in the final objective function (e.g., `convergence_tolerance` = `1.e-6` will result in approximately six digits of accuracy in the final objective function). The `constraint_tolerance` control defines how tightly the constraint functions are satisfied at convergence. The default value is dependent upon the machine precision of the platform in use, but is typically on the order of `1.e-8` for double precision computations. Extremely small values for `constraint_tolerance` may not be attainable. The `output` verbosity setting controls the amount of information generated at each major SQP iteration: the `silent` and `quiet` settings result in only one line of diagnostic output for each major iteration and print the final optimization solution, whereas the `verbose`

and `debug` settings add additional information on the objective function, constraints, and variables at each major iteration.

Concurrency

NPSOL is not a parallel algorithm and cannot directly take advantage of concurrent evaluations. However, if `numerical_gradients` with `method_source dakota` is specified, then the finite difference function evaluations can be performed concurrently (using any of the parallel modes described in the Users Manual [[4] "Adams et al., 2010"]).

An important related observation is the fact that NPSOL uses two different line searches depending on how gradients are computed. For either `analytic_gradients` or `numerical_gradients` with `method_source dakota`, NPSOL is placed in user-supplied gradient mode (NPSOL's "Derivative Level" is set to 3) and it uses a gradient-based line search (the assumption is that user-supplied gradients are inexpensive). On the other hand, if `numerical_gradients` are selected with `method_source vendor`, then NPSOL is computing finite differences internally and it will use a value-based line search (the assumption is that finite differencing on each line search evaluation is too expensive). The ramifications of this are: (1) performance will vary between `method_source dakota` and `method_source vendor` for `numerical_gradients`, and (2) gradient speculation is unnecessary when performing optimization in parallel since the gradient-based line search in user-supplied gradient mode is already load balanced for parallel execution. Therefore, a `speculative` specification will be ignored by NPSOL, and optimization with numerical gradients should select `method_source dakota` for load balanced parallel operation and `method_source vendor` for efficient serial operation.

Linear constraints

Lastly, NPSOL supports specialized handling of linear inequality and equality constraints. By specifying the coefficients and bounds of the linear inequality constraints and the coefficients and targets of the linear equality constraints, this information can be provided to NPSOL at initialization and tracked internally, removing the need for the user to provide the values of the linear constraints on every function evaluation.

verify_level

- Keywords Area
- method
- npsol_sqp
- verify_level

Verify the quality of analytic gradients

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): INTEGER

Description

`verify_level` instructs the algorithm to perform finite difference verifications on analytic gradients. !!!!!!!!!!! used to be: user-supplied gradient components. !!!!!!!!!

function_precision

- Keywords Area
- method
- npsol_sqp
- function_precision

Specify the maximum precision of the analysis code responses

Specification

Alias: none
Argument(s): REAL

Description

The `function_precision` control provides the algorithm with an estimate of the accuracy to which the problem functions can be computed. This is used to prevent the algorithm from trying to distinguish between function values that differ by less than the inherent error in the calculation.

linesearch_tolerance

- Keywords Area
- method
- npsol_sqp
- linesearch_tolerance

Choose how accurately the algorithm will compute the minimum in a line search

Specification

Alias: none
Argument(s): REAL

Description

The `linesearch_tolerance` setting controls the accuracy of the line search. The smaller the value (between 0 and 1), the more accurately the algorithm will attempt to compute a precise minimum along the search direction.

linear_inequality_constraint_matrix

- Keywords Area
- method
- npsol_sqp
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- npsol_sqp
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- npsol_sqp
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- npsol_sqp
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- npsol_sqp
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- npsol_sqp
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- npsol_sqp
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- npsol_sqp
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- npsol_sqp
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- npsol_sqp
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.26 nlssol_sqp

- Keywords Area
- method
- nlssol_sqp

Sequential Quadratic Program for nonlinear least squares

Topics

This keyword is related to the topics:

- sequential_quadratic_programming
- nonlinear_least_squares

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		verify_level	Verify the quality of analytic gradients
	Optional		function_precision	Specify the maximum precision of the analysis code responses
	Optional		linesearch_-tolerance	Choose how accurately the algorithm will compute the minimum in a line search
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint

	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

NLSSOL is available as `nlssol_sqp` and supports unconstrained, bound-constrained, and generally-constrained problems. It exploits the structure of a least squares objective function through the periodic use of Gauss-Newton Hessian approximations to accelerate the SQP algorithm.

Stopping Criteria

The method independent controls for `max_iterations` and `max_function_evaluations` limit the number of major SQP iterations and the number of function evaluations that can be performed during an NPSOL optimization. The `convergence_tolerance` control defines NPSOL's internal optimality tolerance which is used in evaluating if an iterate satisfies the first-order Kuhn-Tucker conditions for a minimum. The magnitude of `convergence_tolerance` approximately specifies the number of significant digits of accuracy desired in the final objective function (e.g., `convergence_tolerance` = `1.e-6` will result in approximately six digits of accuracy in the final objective function). The `constraint_tolerance` control defines how tightly the constraint functions are satisfied at convergence. The default value is dependent upon the machine precision of the platform in use, but is typically on the order of `1.e-8` for double precision computations. Extremely small values for `constraint_tolerance` may not be attainable.

See Also

These keywords may also be of interest:

- npsol_sqp
- nl2sol
- optpp_g_newton

verify_level

- Keywords Area
- method
- nlssol_sqp
- verify_level

Verify the quality of analytic gradients

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): INTEGER

Description

`verify_level` instructs the algorithm to perform finite difference verifications on analytic gradients. !!!!!!!!!!! used to be: user-supplied gradient components. !!!!!!!!!

function_precision

- Keywords Area
- method
- nlssol_sqp
- function_precision

Specify the maximum precision of the analysis code responses

Specification

Alias: none
Argument(s): REAL

Description

The `function_precision` control provides the algorithm with an estimate of the accuracy to which the problem functions can be computed. This is used to prevent the algorithm from trying to distinguish between function values that differ by less than the inherent error in the calculation.

linesearch_tolerance

- Keywords Area
- method
- nlssol_sqp
- linesearch_tolerance

Choose how accurately the algorithm will compute the minimum in a line search

Specification

Alias: none
Argument(s): REAL

Description

The `linesearch_tolerance` setting controls the accuracy of the line search. The smaller the value (between 0 and 1), the more accurately the algorithm will attempt to compute a precise minimum along the search direction.

linear_inequality_constraint_matrix

- Keywords Area
- method
- nlssol_sqp
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- nlssol_sqp
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- nlssol_sqp
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- nlssol_sqp
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- nlssol_sqp
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- nlssol_sqp
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- nlssol_sqp
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- nlssol_sqp
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- nlssol_sqp
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$
$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- nlssol_sqp
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.27 stanford

- Keywords Area
- method
- stanford

Select methods from the Stanford package

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	npsol	Duplicate of `method-npsol-_sqp`
			nlssol	Duplicate of `method-nlssol-_sqp`
	Optional		verify_level	Verify the quality of analytic gradients

	Optional	function_precision	Specify the maximum precision of the analysis code responses
	Optional	linesearch_-tolerance	Choose how accurately the algorithm will compute the minimum in a line search
	Optional	linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities

	Optional	model_pointer	Specify which model block will be used by a method

Description

This is a duplicate. See the method pages here:

- nlssol_sqp
- npsol_sqp

npsol

- Keywords Area
- method
- stanford
- npsol

Duplicate of `method-npsol_sqp`

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

See the page for npsol_sqp

nlssol

- Keywords Area
- method
- stanford
- nlssol

Duplicate of `method-nlssol_sqp`

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

See the page for nlssol_sqp

verify_level

- Keywords Area
- method
- stanford
- verify_level

Verify the quality of analytic gradients

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): INTEGER

Description

`verify_level` instructs the algorithm to perform finite difference verifications on analytic gradients. !!!!!!!!!! used to be: user-supplied gradient components. !!!!!!!!!

function_precision

- Keywords Area
- method
- stanford
- function_precision

Specify the maximum precision of the analysis code responses

Specification

Alias: none
Argument(s): REAL

Description

The `function_precision` control provides the algorithm with an estimate of the accuracy to which the problem functions can be computed. This is used to prevent the algorithm from trying to distinguish between function values that differ by less than the inherent error in the calculation.

linesearch_tolerance

- Keywords Area
- method
- stanford
- linesearch_tolerance

Choose how accurately the algorithm will compute the minimum in a line search

Specification

Alias: none
Argument(s): REAL

Description

The `linesearch_tolerance` setting controls the accuracy of the line search. The smaller the value (between 0 and 1), the more accurately the algorithm will attempt to compute a precise minimum along the search direction.

linear_inequality_constraint_matrix

- Keywords Area
- method
- stanford
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- stanford
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- stanford
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- stanford
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- stanford
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- stanford
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- stanford
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of $0.$ for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- stanford
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- stanford
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$
$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- stanford
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.28 nlpql_sqp

- Keywords Area
- method
- nlpql_sqp

Sequential Quadratic Program

Topics

This keyword is related to the topics:

- package_nlpql
- sequential_quadratic_programming
- local_optimization_methods

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints

	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

NLPQL provides an implementation of sequential quadratic programming through `nlpqp_sqp`. The particular SQP implementation in `nlpql_sqp` uses a variant with distributed and non-monotone line search. Thus, this variant is designed to be more robust in the presence of inaccurate or noisy gradients common in many engineering applications.

The method independent controls for maximum iterations and output verbosity are mapped to NLPQL controls MAXIT and IPRINT, respectively. The maximum number of function evaluations is enforced within the NLPQLPOptimizer class.

linear_inequality_constraint_matrix

- Keywords Area
- method
- nlpql_sqp

- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- nlpql_sqp
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- nlpql_sqp
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- nlpql_sqp
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- nlpql_sqp
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- nlpql_sqp
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- nlpql_sqp
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- nlpql_sqp
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified

- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- nlpql_sqp
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- nlpql_sqp
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.29 optpp_cg

- Keywords Area
- method
- optpp_cg

A conjugate gradient optimization method

Topics

This keyword is related to the topics:

- package_optpp
- local_optimization_methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		max_step	Max change in design point
	Optional		gradient_tolerance	Stopping critiera based on L2 norm of gradient
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional		linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional		linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional		linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional		linear_equality_-targets	Define target values for the linear equality constraints

	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

The conjugate gradient method is an implementation of the Polak-Ribiere approach and handles only unconstrained problems.

See package_optpp for info related to all `optpp` methods.

See Also

These keywords may also be of interest:

- optpp_g_newton
- optpp_pds
- optpp_fd_newton
- optpp_newton
- optpp_g_newton

max_step

- Keywords Area
- method
- optpp_cg
- max_step

Max change in design point

Specification

Alias: none

Argument(s): REAL

Description

The `max_step` control specifies the maximum step that can be taken when computing a change in the current design point (e.g., limiting the Newton step computed from current gradient and Hessian information). It is equivalent to a move limit or a maximum trust region size. The `gradient_tolerance` control defines the threshold value on the L2 norm of the objective function gradient that indicates convergence to an unconstrained minimum (no active constraints). The `gradient_tolerance` control is defined for all gradient-based optimizers.

gradient_tolerance

- Keywords Area
- method
- optpp_cg
- gradient_tolerance

Stopping critiera based on L2 norm of gradient

Specification

Alias: none
Argument(s): REAL

Description

The `gradient_tolerance` control defines the threshold value on the L2 norm of the objective function gradient that indicates convergence to an unconstrained minimum (no active constraints). The `gradient_tolerance` control is defined for all gradient-based optimizers.

linear_inequality_constraint_matrix

- Keywords Area
- method
- optpp_cg
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- optpp_cg
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- optpp_cg
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- optpp_cg
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- optpp_cg
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- optpp_cg
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- optpp_cg
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- optpp_cg
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- optpp_cg
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$
$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- optpp_cg
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.30 optpp_q_newton

- Keywords Area
- method
- optpp_q_newton

Quasi-Newton optimization method

Topics

This keyword is related to the topics:

- package_optpp
- local_optimization_methods

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		search_method	Select a search method for Newton-based optimizers

	Optional	merit_function	Balance goals of reducing objective function and satisfying constraints
	Optional	steplength_to_-boundary	Controls how close to the boundary of the feasible region the algorithm is allowed to move
	Optional	centering_-parameter	Controls how closely the algorithm should follow the "central path"
	Optional	max_step	Max change in design point
	Optional	gradient_tolerance	Stopping critiera based on L2 norm of gradient
	Optional	linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities

	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

This is a Newton method that expects a gradient and computes a low-rank approximation to the Hessian. Each of the Newton-based methods are automatically bound to the appropriate OPT++ algorithm based on the user constraint specification (unconstrained, bound-constrained, or generally-constrained). In the generally-constrained case, the Newton methods use a nonlinear interior-point approach to manage the constraints.

See package_optpp for info related to all `optpp` methods.

search_method

- Keywords Area
- method
- optpp_q_newton
- search_method

Select a search method for Newton-based optimizers

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	value_based_line_-search	Use only function values for line search

			gradient_based_-line_search	Set the search method to use the gradient
			trust_region	Use trust region as the globalization strategy.
			tr_pds	Use direct search as the local search in a trust region method

Description

The `search_method` control is defined for all Newton-based optimizers and is used to select between `trust_region`, `gradient_based_line_search`, and `value_based_line_search` methods. The `gradient_based_line_search` option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, `value_base_line_search` only satisfies the sufficient decrease condition. At each line search iteration, the `gradient_based_line_search` method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the `value_based_line_search` method is preferred to the `gradient_based_line_search` method. Each of these Newton methods additionally supports the `tr_pds` selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a `trust_region` search method is the default for unconstrained problems.

value_based_line_search

- Keywords Area
- method
- optpp_q_newton
- search_method
- value_based_line_search

Use only function values for line search

Specification

Alias: none
Argument(s): none

Description

The `search_method` control is defined for all Newton-based optimizers and is used to select between `trust_region`, `gradient_based_line_search`, and `value_based_line_search` methods. The `gradient_based_line_search` option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, `value_base_line_search` only satisfies the sufficient decrease condition. At each line search iteration, the `gradient_based_line_search` method computes the function and gradient at the trial point. Consequently, given expensive function evaluations,

the value_based_line_search method is preferred to the gradient_based_line_search method. Each of these Newton methods additionally supports the tr_pds selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a trust_region search method is the default for unconstrained problems.

gradient_based_line_search

- Keywords Area
- method
- optpp_q_newton
- search_method
- gradient_based_line_search

Set the search method to use the gradient

Specification

Alias: none
Argument(s): none

Description

The search_method control is defined for all Newton-based optimizers and is used to select between trust_region, gradient_based_line_search, and value_based_line_search methods. The gradient_based_line_search option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, value_base_line_search only satisfies the sufficient decrease condition. At each line search iteration, the gradient_based_line_search method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the value_based_line_search method is preferred to the gradient_based_line_search method. Each of these Newton methods additionally supports the tr_pds selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a trust_region search method is the default for unconstrained problems.

trust_region

- Keywords Area
- method
- optpp_q_newton
- search_method
- trust_region

Use trust region as the globalization strategy.

Specification

Alias: none
Argument(s): none

Description

The trust_region optional group specification can be used to specify the initial size of the trust region (using initial_size) relative to the total variable bounds, the minimum size of the trust region (using minimum_size), the contraction factor for the trust region size (using contraction_factor) used when the surrogate model is performing poorly, and the expansion factor for the trust region size (using expansion_factor) used when the the surrogate model is performing well. Two additional commands are the trust region size contraction threshold (using contract_threshold) and the trust region size expansion threshold (using expand_threshold). These two commands are related to what is called the trust region ratio, which is the actual decrease in the truth model divided by the predicted decrease in the truth model in the current trust region. The command contract_threshold sets the minimum acceptable value for the trust region ratio, i.e., values below this threshold cause the trust region to shrink for the next SBL iteration. The command expand_threshold determines the trust region value above which the trust region will expand for the next SBL iteration.

tr_pds

- Keywords Area
- method
- optpp_q_newton
- search_method
- tr_pds

Use direct search as the local search in a trust region method

Specification

Alias: none
Argument(s): none

Description

The search_method control is defined for all Newton-based optimizers and is used to select between trust_region, gradient_based_line_search, and value_based_line_search methods. The gradient_based_line_search option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, value_base_line_search only satisfies the sufficient decrease condition. At each line search iteration, the gradient_based_line_search method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the value_based_line_search method is preferred to the gradient_based_line_search method. Each of these Newton methods additionally supports the tr_pds selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a trust_region search method is the default for unconstrained problems.

merit_function

- Keywords Area
- method
- optpp_q_newton
- merit_function

Balance goals of reducing objective function and satisfying constraints

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	el_bakry	El-Bakry merit function
			argaez_tapia	The merit function by Tapia and Argaez
			van_shanno	The merit function by Vanderbei and Shanno

Description

A `merit_function` is a function in constrained optimization that attempts to provide joint progress toward reducing the objective function and satisfying the constraints.

el_bakry

- Keywords Area
- method
- optpp_q_newton
- merit_function
- el_bakry

El-Bakry merit function

Specification

Alias: none
Argument(s): none

Description

The "el_bakry" merit function is the L2-norm of the first order optimality conditions for the nonlinear programming problem. The cost per linesearch iteration is n+1 function evaluations. For more information, see [[19] "El-Bakry et al., 1996"].

argaez_tapia

- Keywords Area
- method
- optpp_q_newton
- merit_function
- argaez_tapia

The merit function by Tapia and Argaez

Specification

Alias: none
Argument(s): none

Description

The "argaez_tapia" merit function can be classified as a modified augmented Lagrangian function. The augmented Lagrangian is modified by adding to its penalty term a potential reduction function to handle the perturbed complementarity condition. The cost per linesearch iteration is one function evaluation. For more information, see [[81] "Tapia and Argaez"].

If the function evaluation is expensive or noisy, set the `merit_function` to "argaez_tapia" or "van_shanno".

van_shanno

- Keywords Area
- method
- optpp_q_newton
- merit_function
- van_shanno

The merit function by Vanderbei and Shanno

Specification

Alias: none
Argument(s): none

Description

The "van_shanno" merit function can be classified as a penalty function for the logarithmic barrier formulation of the nonlinear programming problem. The cost per linesearch iteration is one function evaluation. For more information see [[83] "Vanderbei and Shanno, 1999"].

If the function evaluation is expensive or noisy, set the `merit_function` to "argaez_tapia" or "van_shanno".

steplength_to_boundary

- Keywords Area
- method
- optpp_q_newton
- steplength_to_boundary

Controls how close to the boundary of the feasible region the algorithm is allowed to move

Specification

Alias: none
Argument(s): REAL

Description

The `steplength_to_boundary` specification is a parameter (between 0 and 1) that controls how close to the boundary of the feasible region the algorithm is allowed to move. A value of 1 means that the algorithm is allowed to take steps that may reach the boundary of the feasible region. If the user wishes to maintain strict feasibility of the design parameters this value should be less than 1. Default values are .8, .99995, and .95 for the `el_bakry`, `argaez_tapia`, and `van_shanno` merit functions, respectively.

centering_parameter

- Keywords Area
- method
- optpp_q_newton
- centering_parameter

Controls how closely the algorithm should follow the "central path"

Specification

Alias: none
Argument(s): REAL

Description

The `centering_parameter` specification is a parameter (between 0 and 1) that controls how closely the algorithm should follow the "central path". See [[89] "Wright"] for the definition of central path. The larger the value, the more closely the algorithm follows the central path, which results in small steps. A value of 0 indicates that the algorithm will take a pure Newton step. Default values are .2, .2, and .1 for the `el_bakry`, `argaez_tapia`, and `van_shanno` merit functions, respectively.

max_step

- Keywords Area
- method
- optpp_q_newton
- max_step

Max change in design point

Specification

Alias: none
Argument(s): REAL

Description

The `max_step` control specifies the maximum step that can be taken when computing a change in the current design point (e.g., limiting the Newton step computed from current gradient and Hessian information). It is equivalent to a move limit or a maximum trust region size. The `gradient_tolerance` control defines the threshold value on the L2 norm of the objective function gradient that indicates convergence to an unconstrained minimum (no active constraints). The `gradient_tolerance` control is defined for all gradient-based optimizers.

gradient_tolerance

- Keywords Area
- method
- optpp_q_newton
- gradient_tolerance

Stopping critiera based on L2 norm of gradient

Specification

Alias: none
Argument(s): REAL

Description

The `gradient_tolerance` control defines the threshold value on the L2 norm of the objective function gradient that indicates convergence to an unconstrained minimum (no active constraints). The `gradient_tolerance` control is defined for all gradient-based optimizers.

linear_inequality_constraint_matrix

- Keywords Area
- method
- optpp_q_newton
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- optpp_q_newton
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- optpp_q_newton
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- optpp_q_newton
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- optpp_q_newton
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- optpp_q_newton
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- optpp_q_newton
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- optpp_q_newton
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- optpp_q_newton
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- optpp_q_newton
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.31 optpp_fd_newton

- Keywords Area
- method
- optpp_fd_newton

Finite Difference Newton optimization method

Topics

This keyword is related to the topics:

- package_optpp
- local_optimization_methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		search_method	Select a search method for Newton-based optimizers
	Optional		merit_function	Balance goals of reducing objective function and satisfying constraints
	Optional		steplength_to_-boundary	Controls how close to the boundary of the feasible region the algorithm is allowed to move
	Optional		centering_-parameter	Controls how closely the algorithm should follow the ”central path”

	Optional	max_step	Max change in design point
	Optional	gradient_tolerance	Stopping critiera based on L2 norm of gradient
	Optional	linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

This is a Newton method that expects a gradient and computes a finite-difference approximation to the Hessian. Each of the Newton-based methods are automatically bound to the appropriate OPT++ algorithm based on the user constraint specification (unconstrained, bound-constrained, or generally-constrained). In the generally-constrained case, the Newton methods use a nonlinear interior-point approach to manage the constraints.

See package_optpp for info related to all `optpp` methods.

See Also

These keywords may also be of interest:

- optpp_cg
- optpp_g_newton
- optpp_pds
- optpp_newton
- optpp_g_newton

search_method

- Keywords Area
- method
- optpp_fd_newton
- search_method

Select a search method for Newton-based optimizers

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	Description of Group **Group 1**	Dakota Keyword value_based_line_-search	Dakota Keyword Description Use only function values for line search
			gradient_based_-line_search	Set the search method to use the gradient
			trust_region	Use trust region as the globalization strategy.
			tr_pds	Use direct search as the local search in a trust region method

Description

The `search_method` control is defined for all Newton-based optimizers and is used to select between `trust_-region`, `gradient_based_line_search`, and `value_based_line_search` methods. The `gradient_-based_line_search` option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, `value_base_line_search` only satisfies the sufficient decrease condition. At each line search iteration, the `gradient_based_line_search`

method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the value_based_line_search method is preferred to the gradient_based_line_search method. Each of these Newton methods additionally supports the tr_pds selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a trust_region search method is the default for unconstrained problems.

value_based_line_search

- Keywords Area
- method
- optpp_fd_newton
- search_method
- value_based_line_search

Use only function values for line search

Specification

Alias: none
Argument(s): none

Description

The search_method control is defined for all Newton-based optimizers and is used to select between trust_region, gradient_based_line_search, and value_based_line_search methods. The gradient_based_line_search option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, value_base_line_search only satisfies the sufficient decrease condition. At each line search iteration, the gradient_based_line_search method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the value_based_line_search method is preferred to the gradient_based_line_search method. Each of these Newton methods additionally supports the tr_pds selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a trust_region search method is the default for unconstrained problems.

gradient_based_line_search

- Keywords Area
- method
- optpp_fd_newton
- search_method
- gradient_based_line_search

Set the search method to use the gradient

Specification

Alias: none
Argument(s): none

Description

The search_method control is defined for all Newton-based optimizers and is used to select between trust_region, gradient_based_line_search, and value_based_line_search methods. The gradient_based_line_search option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, value_base_line_search only satisfies the sufficient decrease condition. At each line search iteration, the gradient_based_line_search method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the value_based_line_search method is preferred to the gradient_based_line_search method. Each of these Newton methods additionally supports the tr_pds selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a trust_region search method is the default for unconstrained problems.

trust_region

- Keywords Area
- method
- optpp_fd_newton
- search_method
- trust_region

Use trust region as the globalization strategy.

Specification

Alias: none
Argument(s): none

Description

The trust_region optional group specification can be used to specify the initial size of the trust region (using initial_size) relative to the total variable bounds, the minimum size of the trust region (using minimum_size), the contraction factor for the trust region size (using contraction_factor) used when the surrogate model is performing poorly, and the expansion factor for the trust region size (using expansion_factor) used when the the surrogate model is performing well. Two additional commands are the trust region size contraction threshold (using contract_threshold) and the trust region size expansion threshold (using expand_threshold). These two commands are related to what is called the trust region ratio, which is the actual decrease in the truth model divided by the predicted decrease in the truth model in the current trust region. The command contract_threshold sets the minimum acceptable value for the trust region ratio, i.e., values below this threshold cause the trust region to shrink for the next SBL iteration. The command expand_threshold determines the trust region value above which the trust region will expand for the next SBL iteration.

tr_pds

- Keywords Area
- method
- optpp_fd_newton
- search_method
- tr_pds

Use direct search as the local search in a trust region method

Specification

Alias: none
Argument(s): none

Description

The `search_method` control is defined for all Newton-based optimizers and is used to select between `trust_region`, `gradient_based_line_search`, and `value_based_line_search` methods. The `gradient_based_line_search` option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, `value_base_line_search` only satisfies the sufficient decrease condition. At each line search iteration, the `gradient_based_line_search` method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the `value_based_line_search` method is preferred to the `gradient_based_line_search` method. Each of these Newton methods additionally supports the `tr_pds` selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a `trust_region` search method is the default for unconstrained problems.

merit_function

- Keywords Area
- method
- optpp_fd_newton
- merit_function

Balance goals of reducing objective function and satisfying constraints

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	el_bakry	El-Bakry merit function
			argaez_tapia	The merit function by Tapia and Argaez
			van_shanno	The merit function by Vanderbei and Shanno

Description

A `merit_function` is a function in constrained optimization that attempts to provide joint progress toward reducing the objective function and satisfying the constraints.

el_bakry

- Keywords Area
- method
- optpp_fd_newton
- merit_function
- el_bakry

El-Bakry merit function

Specification

Alias: none
Argument(s): none

Description

The "el_bakry" merit function is the L2-norm of the first order optimality conditions for the nonlinear programming problem. The cost per linesearch iteration is n+1 function evaluations. For more information, see [[19] "El-Bakry et al., 1996"].

argaez_tapia

- Keywords Area
- method
- optpp_fd_newton
- merit_function
- argaez_tapia

The merit function by Tapia and Argaez

Specification

Alias: none
Argument(s): none

Description

The "argaez_tapia" merit function can be classified as a modified augmented Lagrangian function. The augmented Lagrangian is modified by adding to its penalty term a potential reduction function to handle the perturbed complementarity condition. The cost per linesearch iteration is one function evaluation. For more information, see [[81] "Tapia and Argaez"].

If the function evaluation is expensive or noisy, set the `merit_function` to "argaez_tapia" or "van_shanno".

van_shanno

- Keywords Area
- method
- optpp_fd_newton
- merit_function
- van_shanno

The merit function by Vanderbei and Shanno

Specification

Alias: none
Argument(s): none

Description

The "van_shanno" merit function can be classified as a penalty function for the logarithmic barrier formulation of the nonlinear programming problem. The cost per linesearch iteration is one function evaluation. For more information see [[83] "Vanderbei and Shanno, 1999"].

If the function evaluation is expensive or noisy, set the `merit_function` to "argaez_tapia" or "van_shanno".

steplength_to_boundary

- Keywords Area
- method
- optpp_fd_newton
- steplength_to_boundary

Controls how close to the boundary of the feasible region the algorithm is allowed to move

Specification

Alias: none
Argument(s): REAL

Description

The `steplength_to_boundary` specification is a parameter (between 0 and 1) that controls how close to the boundary of the feasible region the algorithm is allowed to move. A value of 1 means that the algorithm is allowed to take steps that may reach the boundary of the feasible region. If the user wishes to maintain strict feasibility of the design parameters this value should be less than 1. Default values are .8, .99995, and .95 for the `el_bakry`, `argaez_tapia`, and `van_shanno` merit functions, respectively.

centering_parameter

- Keywords Area
- method
- optpp_fd_newton
- centering_parameter

Controls how closely the algorithm should follow the ”central path”

Specification

Alias: none
Argument(s): REAL

Description

The `centering_parameter` specification is a parameter (between 0 and 1) that controls how closely the algorithm should follow the ”central path”. See [[89] ”Wright”] for the definition of central path. The larger the value, the more closely the algorithm follows the central path, which results in small steps. A value of 0 indicates that the algorithm will take a pure Newton step. Default values are .2, .2, and .1 for the `el_bakry`, `argaez_tapia`, and `van_shanno` merit functions, respectively.

max_step

- Keywords Area
- method
- optpp_fd_newton
- max_step

Max change in design point

Specification

Alias: none
Argument(s): REAL

Description

The `max_step` control specifies the maximum step that can be taken when computing a change in the current design point (e.g., limiting the Newton step computed from current gradient and Hessian information). It is equivalent to a move limit or a maximum trust region size. The `gradient_tolerance` control defines the threshold value on the L2 norm of the objective function gradient that indicates convergence to an unconstrained minimum (no active constraints). The `gradient_tolerance` control is defined for all gradient-based optimizers.

gradient_tolerance

- Keywords Area
- method
- optpp_fd_newton
- gradient_tolerance

Stopping critiera based on L2 norm of gradient

Specification

Alias: none
Argument(s): REAL

Description

The `gradient_tolerance` control defines the threshold value on the L2 norm of the objective function gradient that indicates convergence to an unconstrained minimum (no active constraints). The `gradient_tolerance` control is defined for all gradient-based optimizers.

linear_inequality_constraint_matrix

- Keywords Area
- method
- optpp_fd_newton
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- optpp_fd_newton
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- optpp_fd_newton
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- optpp_fd_newton
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- optpp_fd_newton
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- optpp_fd_newton
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- optpp_fd_newton
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- optpp_fd_newton
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- optpp_fd_newton
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$
$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- optpp_fd_newton
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.32 optpp_g_newton

- Keywords Area
- method
- optpp_g_newton

Newton method based least-squares calbration

Topics

This keyword is related to the topics:

- package_optpp
- local_optimization_methods

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		search_method	Select a search method for Newton-based optimizers

	Optional	merit_function	Balance goals of reducing objective function and satisfying constraints
	Optional	steplength_to_-boundary	Controls how close to the boundary of the feasible region the algorithm is allowed to move
	Optional	centering_-parameter	Controls how closely the algorithm should follow the "central path"
	Optional	max_step	Max change in design point
	Optional	gradient_tolerance	Stopping critiera based on L2 norm of gradient
	Optional	linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities

	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

The Gauss-Newton algorithm is available as `optpp_g_newton` and supports unconstrained, bound-constrained, and generally-constrained problems. When interfaced with the unconstrained, bound-constrained, and nonlinear interior point full-Newton optimizers from the OPT++ library, it provides a Gauss-Newton least squares capability which – on zero-residual test problems – can exhibit quadratic convergence rates near the solution. (Real problems almost never have zero residuals, i.e., perfect fits.)

See package_optpp for info related to all `optpp` methods.

See Also

These keywords may also be of interest:

- optpp_cg
- optpp_pds
- optpp_fd_newton
- optpp_newton
- optpp_g_newton

search_method

- Keywords Area
- method
- optpp_g_newton
- search_method

Select a search method for Newton-based optimizers

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	value_based_line_-search	Use only function values for line search
			gradient_based_-line_search	Set the search method to use the gradient
			trust_region	Use trust region as the globalization strategy.
			tr_pds	Use direct search as the local search in a trust region method

Description

The `search_method` control is defined for all Newton-based optimizers and is used to select between `trust_region`, `gradient_based_line_search`, and `value_based_line_search` methods. The `gradient_based_line_search` option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, `value_base_line_search` only satisfies the sufficient decrease condition. At each line search iteration, the `gradient_based_line_search` method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the `value_based_line_search` method is preferred to the `gradient_based_line_search` method. Each of these Newton methods additionally supports the `tr_pds` selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a `trust_region` search method is the default for unconstrained problems.

value_based_line_search

- Keywords Area
- method
- optpp_g_newton
- search_method
- value_based_line_search

Use only function values for line search

Specification

Alias: none

Argument(s): none

Description

The search_method control is defined for all Newton-based optimizers and is used to select between trust_region, gradient_based_line_search, and value_based_line_search methods. The gradient_based_line_search option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, value_base_line_search only satisfies the sufficient decrease condition. At each line search iteration, the gradient_based_line_search method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the value_based_line_search method is preferred to the gradient_based_line_search method. Each of these Newton methods additionally supports the tr_pds selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a trust_region search method is the default for unconstrained problems.

gradient_based_line_search

- Keywords Area
- method
- optpp_g_newton
- search_method
- gradient_based_line_search

Set the search method to use the gradient

Specification

Alias: none
Argument(s): none

Description

The search_method control is defined for all Newton-based optimizers and is used to select between trust_region, gradient_based_line_search, and value_based_line_search methods. The gradient_based_line_search option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, value_base_line_search only satisfies the sufficient decrease condition. At each line search iteration, the gradient_based_line_search method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the value_based_line_search method is preferred to the gradient_based_line_search method. Each of these Newton methods additionally supports the tr_pds selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a trust_region search method is the default for unconstrained problems.

trust_region

- Keywords Area
- method

- optpp_g_newton
- search_method
- trust_region

Use trust region as the globalization strategy.

Specification

Alias: none
Argument(s): none

Description

The `trust_region` optional group specification can be used to specify the initial size of the trust region (using `initial_size`) relative to the total variable bounds, the minimum size of the trust region (using `minimum_size`), the contraction factor for the trust region size (using `contraction_factor`) used when the surrogate model is performing poorly, and the expansion factor for the trust region size (using `expansion_factor`) used when the the surrogate model is performing well. Two additional commands are the trust region size contraction threshold (using `contract_threshold`) and the trust region size expansion threshold (using `expand_threshold`). These two commands are related to what is called the trust region ratio, which is the actual decrease in the truth model divided by the predicted decrease in the truth model in the current trust region. The command `contract_threshold` sets the minimum acceptable value for the trust region ratio, i.e., values below this threshold cause the trust region to shrink for the next SBL iteration. The command `expand_threshold` determines the trust region value above which the trust region will expand for the next SBL iteration.

tr_pds

- Keywords Area
- method
- optpp_g_newton
- search_method
- tr_pds

Use direct search as the local search in a trust region method

Specification

Alias: none
Argument(s): none

Description

The `search_method` control is defined for all Newton-based optimizers and is used to select between `trust_region`, `gradient_based_line_search`, and `value_based_line_search` methods. The `gradient_based_line_search` option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, `value_base_line_search` only satisfies the sufficient decrease condition. At each line search iteration, the `gradient_based_line_search`

method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the `value_based_line_search` method is preferred to the `gradient_based_line_search` method. Each of these Newton methods additionally supports the `tr_pds` selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a `trust_region` search method is the default for unconstrained problems.

merit_function

- Keywords Area
- method
- optpp_g_newton
- merit_function

Balance goals of reducing objective function and satisfying constraints

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	el_bakry	El-Bakry merit function
			argaez_tapia	The merit function by Tapia and Argaez
			van_shanno	The merit function by Vanderbei and Shanno

Description

A `merit_function` is a function in constrained optimization that attempts to provide joint progress toward reducing the objective function and satisfying the constraints.

el_bakry

- Keywords Area
- method
- optpp_g_newton
- merit_function
- el_bakry

El-Bakry merit function

Specification

Alias: none
Argument(s): none

Description

The "el_bakry" merit function is the L2-norm of the first order optimality conditions for the nonlinear programming problem. The cost per linesearch iteration is n+1 function evaluations. For more information, see [[19] "El-Bakry et al., 1996"].

argaez_tapia

- Keywords Area
- method
- optpp_g_newton
- merit_function
- argaez_tapia

The merit function by Tapia and Argaez

Specification

Alias: none
Argument(s): none

Description

The "argaez_tapia" merit function can be classified as a modified augmented Lagrangian function. The augmented Lagrangian is modified by adding to its penalty term a potential reduction function to handle the perturbed complementarity condition. The cost per linesearch iteration is one function evaluation. For more information, see [[81] "Tapia and Argaez"].

If the function evaluation is expensive or noisy, set the `merit_function` to "argaez_tapia" or "van_shanno".

van_shanno

- Keywords Area
- method
- optpp_g_newton
- merit_function
- van_shanno

The merit function by Vanderbei and Shanno

Specification

Alias: none
Argument(s): none

Description

The "van_shanno" merit function can be classified as a penalty function for the logarithmic barrier formulation of the nonlinear programming problem. The cost per linesearch iteration is one function evaluation. For more information see [[83] "Vanderbei and Shanno, 1999"].

If the function evaluation is expensive or noisy, set the `merit_function` to "argaez_tapia" or "van_shanno".

steplength_to_boundary

- Keywords Area
- method
- optpp_g_newton
- steplength_to_boundary

Controls how close to the boundary of the feasible region the algorithm is allowed to move

Specification

Alias: none
Argument(s): REAL

Description

The `steplength_to_boundary` specification is a parameter (between 0 and 1) that controls how close to the boundary of the feasible region the algorithm is allowed to move. A value of 1 means that the algorithm is allowed to take steps that may reach the boundary of the feasible region. If the user wishes to maintain strict feasibility of the design parameters this value should be less than 1. Default values are .8, .99995, and .95 for the `el_bakry`, `argaez_tapia`, and `van_shanno` merit functions, respectively.

centering_parameter

- Keywords Area
- method
- optpp_g_newton
- centering_parameter

Controls how closely the algorithm should follow the "central path"

Specification

Alias: none
Argument(s): REAL

Description

The `centering_parameter` specification is a parameter (between 0 and 1) that controls how closely the algorithm should follow the "central path". See [[89] "Wright"] for the definition of central path. The larger the value, the more closely the algorithm follows the central path, which results in small steps. A value of 0 indicates that the algorithm will take a pure Newton step. Default values are .2, .2, and .1 for the `el_bakry`, `argaez_tapia`, and `van_shanno` merit functions, respectively.

max_step

- Keywords Area
- method
- optpp_g_newton
- max_step

Max change in design point

Specification

Alias: none
Argument(s): REAL

Description

The `max_step` control specifies the maximum step that can be taken when computing a change in the current design point (e.g., limiting the Newton step computed from current gradient and Hessian information). It is equivalent to a move limit or a maximum trust region size. The `gradient_tolerance` control defines the threshold value on the L2 norm of the objective function gradient that indicates convergence to an unconstrained minimum (no active constraints). The `gradient_tolerance` control is defined for all gradient-based optimizers.

gradient_tolerance

- Keywords Area
- method
- optpp_g_newton
- gradient_tolerance

Stopping critiera based on L2 norm of gradient

Specification

Alias: none
Argument(s): REAL

Description

The `gradient_tolerance` control defines the threshold value on the L2 norm of the objective function gradient that indicates convergence to an unconstrained minimum (no active constraints). The `gradient_tolerance` control is defined for all gradient-based optimizers.

linear_inequality_constraint_matrix

- Keywords Area
- method
- optpp_g_newton
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- optpp_g_newton
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- optpp_g_newton
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- optpp_g_newton
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- optpp_g_newton
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- optpp_g_newton
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- optpp_g_newton
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- optpp_g_newton
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- optpp_g_newton
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- optpp_g_newton
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.33 optpp_newton

- Keywords Area
- method
- optpp_newton

Newton method based optimization

Topics

This keyword is related to the topics:

- package_optpp
- local_optimization_methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		search_method	Select a search method for Newton-based optimizers
	Optional		merit_function	Balance goals of reducing objective function and satisfying constraints
	Optional		steplength_to_-boundary	Controls how close to the boundary of the feasible region the algorithm is allowed to move
	Optional		centering_-parameter	Controls how closely the algorithm should follow the "central path"

	Optional	max_step	Max change in design point
	Optional	gradient_tolerance	Stopping critiera based on L2 norm of gradient
	Optional	linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

This is a full Newton method that expects a gradient and a Hessian. Each of the Newton-based methods are automatically bound to the appropriate OPT++ algorithm based on the user constraint specification (unconstrained, bound-constrained, or generally-constrained). In the generally-constrained case, the Newton methods use a nonlinear interior-point approach to manage the constraints.

See package_optpp for info related to all `optpp` methods.

See Also

These keywords may also be of interest:

- optpp_cg
- optpp_g_newton
- optpp_pds
- optpp_fd_newton
- optpp_g_newton

search_method

- Keywords Area
- method
- optpp_newton
- search_method

Select a search method for Newton-based optimizers

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** value_based_line_-search	**Dakota Keyword Description** Use only function values for line search
			gradient_based_-line_search	Set the search method to use the gradient

			trust_region	Use trust region as the globalization strategy.
			tr_pds	Use direct search as the local search in a trust region method

Description

The `search_method` control is defined for all Newton-based optimizers and is used to select between `trust_region`, `gradient_based_line_search`, and `value_based_line_search` methods. The `gradient_based_line_search` option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, `value_base_line_search` only satisfies the sufficient decrease condition. At each line search iteration, the `gradient_based_line_search` method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the `value_based_line_search` method is preferred to the `gradient_based_line_search` method. Each of these Newton methods additionally supports the `tr_pds` selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a `trust_region` search method is the default for unconstrained problems.

value_based_line_search

- Keywords Area
- method
- optpp_newton
- search_method
- value_based_line_search

Use only function values for line search

Specification

Alias: none
Argument(s): none

Description

The `search_method` control is defined for all Newton-based optimizers and is used to select between `trust_region`, `gradient_based_line_search`, and `value_based_line_search` methods. The `gradient_based_line_search` option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, `value_base_line_search` only satisfies the sufficient decrease condition. At each line search iteration, the `gradient_based_line_search` method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the `value_based_line_search` method is preferred to the `gradient_based_line_search` method. Each of these Newton methods additionally supports the `tr_pds` selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default

for bound-constrained and generally-constrained problems, and use of a `trust_region` search method is the default for unconstrained problems.

gradient_based_line_search

- Keywords Area
- method
- optpp_newton
- search_method
- gradient_based_line_search

Set the search method to use the gradient

Specification

Alias: none
Argument(s): none

Description

The `search_method` control is defined for all Newton-based optimizers and is used to select between `trust_region`, `gradient_based_line_search`, and `value_based_line_search` methods. The `gradient_based_line_search` option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, `value_base_line_search` only satisfies the sufficient decrease condition. At each line search iteration, the `gradient_based_line_search` method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the `value_based_line_search` method is preferred to the `gradient_based_line_search` method. Each of these Newton methods additionally supports the `tr_pds` selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a `trust_region` search method is the default for unconstrained problems.

trust_region

- Keywords Area
- method
- optpp_newton
- search_method
- trust_region

Use trust region as the globalization strategy.

Specification

Alias: none
Argument(s): none

Description

The trust_region optional group specification can be used to specify the initial size of the trust region (using initial_size) relative to the total variable bounds, the minimum size of the trust region (using minimum_size), the contraction factor for the trust region size (using contraction_factor) used when the surrogate model is performing poorly, and the expansion factor for the trust region size (using expansion_factor) used when the the surrogate model is performing well. Two additional commands are the trust region size contraction threshold (using contract_threshold) and the trust region size expansion threshold (using expand_threshold). These two commands are related to what is called the trust region ratio, which is the actual decrease in the truth model divided by the predicted decrease in the truth model in the current trust region. The command contract_threshold sets the minimum acceptable value for the trust region ratio, i.e., values below this threshold cause the trust region to shrink for the next SBL iteration. The command expand_threshold determines the trust region value above which the trust region will expand for the next SBL iteration.

tr_pds

- Keywords Area
- method
- optpp_newton
- search_method
- tr_pds

Use direct search as the local search in a trust region method

Specification

Alias: none
Argument(s): none

Description

The search_method control is defined for all Newton-based optimizers and is used to select between trust_region, gradient_based_line_search, and value_based_line_search methods. The gradient_based_line_search option uses the line search method proposed by [[62] "More and Thuente, 1994"]. This option satisfies sufficient decrease and curvature conditions; whereas, value_base_line_search only satisfies the sufficient decrease condition. At each line search iteration, the gradient_based_line_search method computes the function and gradient at the trial point. Consequently, given expensive function evaluations, the value_based_line_search method is preferred to the gradient_based_line_search method. Each of these Newton methods additionally supports the tr_pds selection for unconstrained problems. This option performs a robust trust region search using pattern search techniques. Use of a line search is the default for bound-constrained and generally-constrained problems, and use of a trust_region search method is the default for unconstrained problems.

merit_function

- Keywords Area
- method

- optpp_newton
- merit_function

Balance goals of reducing objective function and satisfying constraints

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	el_bakry	El-Bakry merit function
			argaez_tapia	The merit function by Tapia and Argaez
			van_shanno	The merit function by Vanderbei and Shanno

Description

A `merit_function` is a function in constrained optimization that attempts to provide joint progress toward reducing the objective function and satisfying the constraints.

el_bakry

- Keywords Area
- method
- optpp_newton
- merit_function
- el_bakry

El-Bakry merit function

Specification

Alias: none
Argument(s): none

Description

The "el_bakry" merit function is the L2-norm of the first order optimality conditions for the nonlinear programming problem. The cost per linesearch iteration is n+1 function evaluations. For more information, see [[19] "El-Bakry et al., 1996"].

argaez_tapia

- Keywords Area
- method
- optpp_newton
- merit_function
- argaez_tapia

The merit function by Tapia and Argaez

Specification

Alias: none
Argument(s): none

Description

The "argaez_tapia" merit function can be classified as a modified augmented Lagrangian function. The augmented Lagrangian is modified by adding to its penalty term a potential reduction function to handle the perturbed complementarity condition. The cost per linesearch iteration is one function evaluation. For more information, see [[81] "Tapia and Argaez"].

If the function evaluation is expensive or noisy, set the `merit_function` to "argaez_tapia" or "van_shanno".

van_shanno

- Keywords Area
- method
- optpp_newton
- merit_function
- van_shanno

The merit function by Vanderbei and Shanno

Specification

Alias: none
Argument(s): none

Description

The "van_shanno" merit function can be classified as a penalty function for the logarithmic barrier formulation of the nonlinear programming problem. The cost per linesearch iteration is one function evaluation. For more information see [[83] "Vanderbei and Shanno, 1999"].

If the function evaluation is expensive or noisy, set the `merit_function` to "argaez_tapia" or "van_shanno".

steplength_to_boundary

- Keywords Area
- method
- optpp_newton
- steplength_to_boundary

Controls how close to the boundary of the feasible region the algorithm is allowed to move

Specification

Alias: none
Argument(s): REAL

Description

The `steplength_to_boundary` specification is a parameter (between 0 and 1) that controls how close to the boundary of the feasible region the algorithm is allowed to move. A value of 1 means that the algorithm is allowed to take steps that may reach the boundary of the feasible region. If the user wishes to maintain strict feasibility of the design parameters this value should be less than 1. Default values are .8, .99995, and .95 for the `el_bakry`, `argaez_tapia`, and `van_shanno` merit functions, respectively.

centering_parameter

- Keywords Area
- method
- optpp_newton
- centering_parameter

Controls how closely the algorithm should follow the ”central path”

Specification

Alias: none
Argument(s): REAL

Description

The `centering_parameter` specification is a parameter (between 0 and 1) that controls how closely the algorithm should follow the ”central path”. See [[89] ”Wright”] for the definition of central path. The larger the value, the more closely the algorithm follows the central path, which results in small steps. A value of 0 indicates that the algorithm will take a pure Newton step. Default values are .2, .2, and .1 for the `el_bakry`, `argaez_tapia`, and `van_shanno` merit functions, respectively.

max_step

- Keywords Area
- method
- optpp_newton
- max_step

Max change in design point

Specification

Alias: none
Argument(s): REAL

Description

The `max_step` control specifies the maximum step that can be taken when computing a change in the current design point (e.g., limiting the Newton step computed from current gradient and Hessian information). It is equivalent to a move limit or a maximum trust region size. The `gradient_tolerance` control defines the threshold value on the L2 norm of the objective function gradient that indicates convergence to an unconstrained minimum (no active constraints). The `gradient_tolerance` control is defined for all gradient-based optimizers.

gradient_tolerance

- Keywords Area
- method
- optpp_newton
- gradient_tolerance

Stopping critiera based on L2 norm of gradient

Specification

Alias: none
Argument(s): REAL

Description

The `gradient_tolerance` control defines the threshold value on the L2 norm of the objective function gradient that indicates convergence to an unconstrained minimum (no active constraints). The `gradient_tolerance` control is defined for all gradient-based optimizers.

linear_inequality_constraint_matrix

- Keywords Area
- method
- optpp_newton
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- optpp_newton
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- optpp_newton
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- optpp_newton
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- optpp_newton
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- optpp_newton
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- optpp_newton
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- optpp_newton
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- optpp_newton
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- optpp_newton
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.34 optpp_pds

- Keywords Area
- method
- optpp_pds

Simplex-based derivative free optimization method

Topics

This keyword is related to the topics:

- package_optpp

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		search_scheme_-size	Number of points to be used in the direct search template
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional		linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional		linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional		linear_inequality_-scales	Define the characteristic values to scale linear inequalities

	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

The direct search algorithm, PDS (parallel direct search method), supports bound constraints.

The PDS method can directly exploit asynchronous evaluations; however, this capability has not yet been implemented in Dakota.

See package_optpp for info related to all `optpp` methods.

See Also

These keywords may also be of interest:

- optpp_cg
- optpp_g_newton
- optpp_fd_newton
- optpp_newton
- optpp_g_newton

search_scheme_size

- Keywords Area
- method
- optpp_pds
- search_scheme_size

Number of points to be used in the direct search template

Specification

Alias: none
Argument(s): INTEGER

Description

The `search_scheme_size` is defined for the PDS method to specify the number of points to be used in the direct search template.

linear_inequality_constraint_matrix

- Keywords Area
- method
- optpp_pds
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- optpp_pds
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- optpp_pds
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- optpp_pds
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- optpp_pds
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- optpp_pds
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- optpp_pds
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- optpp_pds
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- optpp_pds
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- optpp_pds
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.35 asynch_pattern_search

- Keywords Area
- method
- asynch_pattern_search

Pattern search, derivative free optimization method

Topics

This keyword is related to the topics:

- package_hopspack
- global_optimization_methods

Specification

Alias: coliny_apps

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_delta	Initial step size for non-gradient based optimizers
	Optional		contraction_factor	Amount by which step length is rescaled
	Optional		threshold_delta	Stopping criteria based on step length or pattern size
	Optional		solution_target	Stopping criteria based on objective function value
	Optional		synchronization	Select how Dakota schedules function evaluations
	Optional		merit_function	Balance goals of reducing objective function and satisfying constraints

	Optional	constraint_penalty	Multiplier for the penalty function
	Optional	smoothing_factor	Smoothing value for smoothed penalty functions
	Optional	linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

The asynchronous parallel pattern search (APPS) algorithm [[37] "Gray and Kolda, 2006"] is a fully asynchronous pattern search technique in that the search along each offset direction continues without waiting for searches along other directions to finish.

Currently, APPS only supports coordinate bases with a total of *2n* function evaluations in the pattern, and these patterns may only contract.

Concurrency

APPS exploits parallelism through the use of Dakota's concurrent function evaluations. The variant of the algorithm that is currently exposed, however, limits the amount of concurrency that can be exploited. In particular, APPS can leverage an evaluation concurrency level of at most twice the number of variables. More options that allow for greater evaluation concurrency may be exposed in future releases.

Algorithm Behavior

- `initial_delta`: the initial step length, must be positive
- `threshold_delta`: step length used to determine convergence, must be greater than or equal to 4.4e-16
- `contraction_factor`: amount by which step length is rescaled after unsuccesful iterates, must be strictly between 0 and 1

Merit Functions

APPS solves nonlinearly constrained problems by solving a sequence of linearly constrained merit function-base subproblems. There are several exact and smoothed exact penalty functions that can be specified with the `merit_function` control. The options are as follows:

- `merit_max`: based on ℓ_∞ norm
- `merit_max_smooth`: based on smoothed ℓ_∞ norm
- `merit1`: based on ℓ_1 norm
- `merit1_smooth`: based on smoothed ℓ_1 norm
- `merit2`: based on ℓ_2 norm
- `merit2_smooth`: based on smoothed ℓ_2 norm
- `merit2_squared`: based on ℓ_2^2 norm

The user can also specify the following to affect the merit functions:

- `constraint_penalty`
- `smoothing_parameter`

Method Independent Controls

The only method independent controls that are currently mapped to APPS are:

- max_function_evaluations
- constraint_tolerance
- output

Note that while APPS treats the constraint tolerance separately for linear and nonlinear constraints, we apply the same value to both if the user specifies `constraint_tolerance`.

The APPS internal display level is mapped to the Dakota `output` settings as follows:

- `debug`: display final solution, all input parameters, variable and constraint info, trial points, search directions, and execution details
- `verbose`: display final solution, all input parameters, variable and constraint info, and trial points
- `normal`: display final solution, all input parameters, variable and constraint summaries, and new best points
- `quiet`: display final solution and all input parameters
- `silent`: display final solution

initial_delta

- Keywords Area
- method
- asynch_pattern_search
- initial_delta

Initial step size for non-gradient based optimizers

Specification

Alias: none

Argument(s): REAL

Description

If `initial_delta` is supplied by the user, it will be applied in an absolute sense in all coordinate directions. APPS documentation advocates choosing `initial_delta` to be the approximate distance from the initial point to the solution. If this is unknown, it is advisable to err on the side of choosing an `initial_delta` that is too large or to not specify it. In the latter case, APPS will take a full step to the boundary in each direction. Relative application of `initial_delta` is not available unless the user scales the problem accordingly.

contraction_factor

- Keywords Area
- method
- asynch_pattern_search
- contraction_factor

Amount by which step length is rescaled

Specification

Alias: none
Argument(s): REAL

Description

For pattern search methods, `contraction_factor` specifies the amount by which step length is rescaled after unsuccesful iterates, must be strictly between 0 and 1.

For methods that can expand the step length, the expansion is 1/ `contraction_factor`

threshold_delta

- Keywords Area
- method
- asynch_pattern_search
- threshold_delta

Stopping criteria based on step length or pattern size

Specification

Alias: none
Argument(s): REAL

Description

`threshold_delta` is the step length or pattern size used to determine convergence.

solution_target

- Keywords Area
- method
- asynch_pattern_search
- solution_target

Stopping criteria based on objective function value

Topics

This keyword is related to the topics:

- problem

Specification

Alias: solution_accuracy
Argument(s): REAL

Description

`solution_target` is a termination criterion. The algorithm will terminate when the function value falls below solution_target.

synchronization

- Keywords Area
- method
- asynch_pattern_search
- synchronization

Select how Dakota schedules function evaluations

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** blocking	**Dakota Keyword Description** The `blocking` option causes Dakota to run the algorithm in synchronous mode
			nonblocking	The `nonblocking` option causes Dakota to run the algorithm in asynchronous mode

Description

The `synchronization` specification can be used to specify the use of either `blocking` or `nonblocking` schedulers.

blocking

- Keywords Area
- method

- asynch_pattern_search
- synchronization
- blocking

The `blocking` option causes Dakota to run the algorithm in synchronous mode

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The `blocking` option causes Dakota to run the algorithm in synchronous mode.

nonblocking

- Keywords Area
- method
- asynch_pattern_search
- synchronization
- nonblocking

The `nonblocking` option causes Dakota to run the algorithm in asynchronous mode

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The `nonblocking` option causes Dakota to run the algorithm in asynchronous mode.
The `nonblocking` option is not available when Dakota is used in message-passing mode.

merit_function

- Keywords Area
- method
- asynch_pattern_search
- merit_function

Balance goals of reducing objective function and satisfying constraints

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword merit_max	Dakota Keyword Description Nonsmoothed merit function
	Required*(Choose One)*	**Group 1**	merit_max_smooth	Smoothed merit function
			merit1	Nonsmoothed merit function
			merit1_smooth	Smoothed merit function
			merit2	Nonsmoothed merit function
			merit2_smooth	Smoothed merit function
			merit2_squared	Nonsmoothed merit function

Description

A `merit_function` is a function in constrained optimization that attempts to provide joint progress toward reducing the objective function and satisfying the constraints.

merit_max

- Keywords Area
- method
- asynch_pattern_search
- merit_function
- merit_max

Nonsmoothed merit function

Specification

Alias: none
Argument(s): none

Description

APPS solves nonlinearly constrained problems by solving a sequence of linearly constrained merit function-base subproblems. There are several exact and smoothed exact penalty functions.

`merit_max`: based on ℓ_∞ norm

merit_max_smooth

- Keywords Area
- method
- asynch_pattern_search
- merit_function
- merit_max_smooth

Smoothed merit function

Specification

Alias: none
Argument(s): none

Description

APPS solves nonlinearly constrained problems by solving a sequence of linearly constrained merit function-base subproblems. There are several exact and smoothed exact penalty functions.

`merit_max_smooth`: based on smoothed ℓ_∞ norm

merit1

- Keywords Area
- method
- asynch_pattern_search
- merit_function
- merit1

Nonsmoothed merit function

Specification

Alias: none
Argument(s): none

Description

APPS solves nonlinearly constrained problems by solving a sequence of linearly constrained merit function-base subproblems. There are several exact and smoothed exact penalty functions.

- `merit1:` based on ℓ_1 norm

merit1_smooth

- Keywords Area
- method
- asynch_pattern_search
- merit_function
- merit1_smooth

Smoothed merit function

Specification

Alias: none
Argument(s): none

Description

APPS solves nonlinearly constrained problems by solving a sequence of linearly constrained merit function-base subproblems. There are several exact and smoothed exact penalty functions.
`merit1_smooth:` based on smoothed ℓ_1 norm

merit2

- Keywords Area
- method
- asynch_pattern_search
- merit_function
- merit2

Nonsmoothed merit function

Specification

Alias: none
Argument(s): none

Description

APPS solves nonlinearly constrained problems by solving a sequence of linearly constrained merit function-base subproblems. There are several exact and smoothed exact penalty functions.
`merit2:` based on ℓ_2 norm

merit2_smooth

- Keywords Area
- method
- asynch_pattern_search
- merit_function
- merit2_smooth

Smoothed merit function

Specification

Alias: none
Argument(s): none

Description

APPS solves nonlinearly constrained problems by solving a sequence of linearly constrained merit function-base subproblems. There are several exact and smoothed exact penalty functions.

`merit2_smooth`: based on smoothed ℓ_2 norm

merit2_squared

- Keywords Area
- method
- asynch_pattern_search
- merit_function
- merit2_squared

Nonsmoothed merit function

Specification

Alias: none
Argument(s): none

Description

APPS solves nonlinearly constrained problems by solving a sequence of linearly constrained merit function-base subproblems. There are several exact and smoothed exact penalty functions.

`merit2_squared`: based on ℓ_2^2 norm

constraint_penalty

- Keywords Area
- method
- asynch_pattern_search
- constraint_penalty

Multiplier for the penalty function

Specification

Alias: none
Argument(s): REAL

Description

- `constraint_penalty`: the initial multiplier for the penalty function, must be nonnegative

Most SCOLIB optimizers treat constraints with a simple penalty scheme that adds `constraint_penalty` times the sum of squares of the constraint violations to the objective function. Specific exceptions to this method for handling constraint violations are noted below. (The default value of `constraint_penalty` is 1000.0, except for methods that dynamically adapt their constraint penalty, for which the default value is 1.0.)

smoothing_factor

- Keywords Area
- method
- asynch_pattern_search
- smoothing_factor

Smoothing value for smoothed penalty functions

Specification

Alias: none
Argument(s): REAL

Description

- `smoothing_parameter`: initial smoothing value for smoothed penalty functions, must be between 0 and 1 (inclusive)

linear_inequality_constraint_matrix

- Keywords Area
- method
- asynch_pattern_search
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- asynch_pattern_search
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- asynch_pattern_search
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- asynch_pattern_search
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- asynch_pattern_search
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- asynch_pattern_search
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- asynch_pattern_search
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- asynch_pattern_search
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- asynch_pattern_search
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- asynch_pattern_search
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.36 mesh_adaptive_search

- Keywords Area
- method
- mesh_adaptive_search

Mesh Adaptive Direct Search Algorithm

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		function_precision	Specify the maximum precision of the analysis code responses
	Optional		seed	Specify the seed of the random number generator
	Optional		history_file	This is the name of a file where NOMAD will write its own output.
	Optional		display_format	Define output format.
	Optional		variable_-neighborhood_-search	Set parameter to escape local minima
	Optional		display_all_-evaluations	Show all evaluations
	Optional		linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional		linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint

	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

The mesh adaptive direct search algorithm [[8] "Audet, Le Digabel, and Tribes, 2009"] is a generalized pattern search in which the set of points evaluated becomes increasingly dense, leading to good convergence properties. It is now made available in Dakota through the NOMAD software [[2] "Abramson, Audet, Couture, Dennis, Le Digabel, and Tribes"]. It can handle unconstrained problems as well as those with bound constraints and general nonlinear constraints. Of particular note, it can handle both continuous and discrete parameters. NOMAD is available to the public under the GNU LGPL and the source code is included with Dakota. NOMAD-specific software documentation is available from `http://www.gerad.ca/nomad`.

function_precision

- Keywords Area
- method
- mesh_adaptive_search
- function_precision

Specify the maximum precision of the analysis code responses

Specification

Alias: none
Argument(s): REAL

Description

The `function_precision` control provides the algorithm with an estimate of the accuracy to which the problem functions can be computed. This is used to prevent the algorithm from trying to distinguish between function values that differ by less than the inherent error in the calculation.

seed

- Keywords Area
- method
- mesh_adaptive_search
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

history_file

- Keywords Area
- method
- mesh_adaptive_search
- history_file

This is the name of a file where NOMAD will write its own output.

Specification

Alias: none
Argument(s): STRING

Description

`history_file`: This is the name of a file where NOMAD will write its own output.

display_format

- Keywords Area
- method
- mesh_adaptive_search
- display_format

Define output format.

Specification

Alias: none
Argument(s): STRING

Description

Define output format.

variable_neighborhood_search

- Keywords Area
- method
- mesh_adaptive_search
- variable_neighborhood_search

Set parameter to escape local minima

Specification

Alias: none
Argument(s): REAL

Description

`vns`: This parameter guides a variable neighborhood. It roughly corresponds to how aggressive NOMAD is in trying to move away from local minima.

display_all_evaluations

- Keywords Area
- method
- mesh_adaptive_search
- display_all_evaluations

Show all evaluations

Specification

Alias: none
Argument(s): none

Description

`display_all_evaluations`: If this is set, then NOMAD will print out its own record of evaluations completed.

linear_inequality_constraint_matrix

- Keywords Area
- method
- mesh_adaptive_search
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- mesh_adaptive_search
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- mesh_adaptive_search
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- mesh_adaptive_search
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- mesh_adaptive_search
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- mesh_adaptive_search
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- mesh_adaptive_search
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- mesh_adaptive_search
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- mesh_adaptive_search
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none

Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- mesh_adaptive_search
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.37 moga

- Keywords Area
- method
- moga

Multi-objective Genetic Algorithm (a.k.a Evolutionary Algorithm)

Topics

This keyword is related to the topics:

- package_jega

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		fitness_type	Select the fitness type for JEGA methods
	Optional		replacement_type	Select a replacement type for JEGA methods
	Optional		niching_type	Specify the type of niching pressure
	Optional		convergence_type	Select the convergence type for JEGA methods
	Optional		postprocessor_type	Post process the final solution from `moga`
	Optional		population_size	Set the initial population size in JEGA methods
	Optional		log_file	Specify the name of a log file
	Optional		print_each_pop	Print every population to a population file

	Optional	initialization_type	Specify how to initialize the population
	Optional	crossover_type	Select a crossover type for JEGA methods
	Optional	mutation_type	Select a mutation type for JEGA methods
	Optional	seed	Specify the seed of the random number generator
	Optional	linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities
	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities

	Optional	model_pointer	Specify which model block will be used by a method

Description

moga stands for Multi-objective Genetic Algorithm, which is a global optimization method that does Pareto optimization for multiple objectives. It supports general constraints and a mixture of real and discrete variables.

Constraints

moga can utilize linear constraints using the keywords: -∗ linear_inequality_constraint_matrix -∗ linear_inequality_lower_bounds -∗ linear_inequality_upper_bounds -∗ linear_inequality_scale_types -∗ linear_inequality_scales -∗ linear_equality_constraint_matrix -∗ linear_equality_targets -∗ linear_equality_scale_types -∗ linear_equality_scales

Configuration

The genetic algorithm configurations are:

1. fitness
2. replacement
3. niching
4. convergence
5. postprocessor
6. initialization
7. crossover
8. mutation
9. population size

The steps followed by the algorithm are listed below. The configurations will effect how the algorithm completes each step.

Stopping Criteria

The moga method respects the max_iterations and max_function_evaluations method independent controls to provide integer limits for the maximum number of generations and function evaluations, respectively.

The algorithm also stops when convergence is reached. This involves repeated assessment of the algorithm's progress in solving the problem, until some criterion is met.

The specification for convergence in a moga can either be metric_tracker or can be omitted all together. If omitted, no convergence algorithm will be used and the algorithm will rely on stopping criteria only.

Outputs

The moga method respects the output method independent control to vary the amount of information presented to the user during execution.

The final results are written to the Dakota tabular output. Additional information is also available - see the log_file and print_each_pop keywords.

Note that moga and SOGA create additional output files during execution. "finaldata.dat" is a file that holds the final set of Pareto optimal solutions after any post-processing is complete. "discards.dat" holds solutions that were discarded from the population during the course of evolution.

It can often be useful to plot objective function values from these files to visually see the Pareto front and ensure that finaldata.dat solutions dominate discards.dat solutions. The solutions are written to these output files in the format "Input1...InputN..Output1...OutputM".

Important Notes

The pool of potential members is the current population and the current set of offspring.

Choice of fitness assessors is strongly related to the type of replacement algorithm being used and can have a profound effect on the solutions selected for the next generation.

If using the fitness types `layer_rank` or `domination_count`, it is strongly recommended that you use the `replacement_type` `below_limit` (although the roulette wheel selectors can also be used).

The functionality of the domination_count selector of JEGA v1.0 can now be achieved using the `domination_count` fitness type and `below_limit` replacement type.

Theory

The basic steps of the `moga` algorithm are as follows:

1. Initialize the population
2. Evaluate the population (calculate the values of the objective function and constraints for each population member)
3. Loop until converged, or stopping criteria reached
   (a) Perform crossover
   (b) Perform mutation
   (c) Evaluate the new population
   (d) Assess the fitness of each member in the population
   (e) Replace the population with members selected to continue in the next generation
   (f) Apply niche pressure to the population
   (g) Test for convergence
4. Perform post processing

If moga is used in a hybrid optimization method (which requires one optimal solution from each individual optimization method to be passed to the subsequent optimization method as its starting point), the solution in the Pareto set closest to the "utopia" point is given as the best solution. This solution is also reported in the Dakota output.

This "best" solution in the Pareto set has minimum distance from the utopia point. The utopia point is defined as the point of extreme (best) values for each objective function. For example, if the Pareto front is bounded by (1,100) and (90,2), then (1,2) is the utopia point. There will be a point in the Pareto set that has minimum L2-norm distance to this point, for example (10,10) may be such a point.

If moga is used in a method which may require passing multiple solutions to the next level (such as the `surrogate_based_global` method or `hybrid` methods), the `orthogonal_distance` postprocessor type may be used to specify the distances between each solution value to winnow down the solutions in the full Pareto front to a subset which will be passed to the next iteration.

See Also

These keywords may also be of interest:

- soga

fitness_type

- Keywords Area
- method
- moga
- fitness_type

Select the fitness type for JEGA methods

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** layer_rank	**Dakota Keyword Description** Assign each member to a layer, based on domination the rank based on layers
			domination_count	Rank each member by the number of members that dominate it

Description

The two JEGA methods use different fitness types, which are described on their respective pages.

layer_rank

- Keywords Area
- method
- moga
- fitness_type
- layer_rank

Assign each member to a layer, based on domination the rank based on layers

Specification

Alias: none
Argument(s): none

Description

The `fitness_type: layer_rank` has been specifically designed to avoid problems with aggregating and scaling objective function values and transforming them into a single objective.

The `layer_rank` fitness assessor works by assigning all non-dominated designs a layer of 0, then from what remains, assigning all the non-dominated a layer of 1, and so on until all designs have been assigned a layer. The values are negated to follow the higher-is-better fitness convention.

Use of the `below_limit` selector with the `layer_rank` fitness assessor has the effect of keeping all those designs whose layer is below a certain threshold again subject to the shrinkage limit.

domination_count

- Keywords Area
- method
- moga
- fitness_type
- domination_count

Rank each member by the number of members that dominate it

Specification

Alias: none

Argument(s): none

Description

The `fitness_type: domination_count` has been specifically designed to avoid problems with aggregating and scaling objective function values and transforming them into a single objective.

Instead, the `domination_count` fitness assessor works by ordering population members by the negative of the number of designs that dominate them. The values are negated in keeping with the convention that higher fitness is better.

The `layer_rank` fitness assessor works by assigning all non-dominated designs a layer of 0, then from what remains, assigning all the non-dominated a layer of 1, and so on until all designs have been assigned a layer. Again, the values are negated for the higher-is-better fitness convention.

Use of the `below_limit` selector with the `domination_count` fitness assessor has the effect of keeping all designs that are dominated by fewer then a limiting number of other designs subject to the shrinkage limit.

Using it with the `layer_rank` fitness assessor has the effect of keeping all those designs whose layer is below a certain threshold again subject to the shrinkage limit.

replacement_type

- Keywords Area
- method
- moga
- replacement_type

Select a replacement type for JEGA methods

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	Description of Group **Group 1**	Dakota Keyword elitist	Dakota Keyword Description Use the best designs to form a new population
			roulette_wheel	Replace population
			unique_roulette_-wheel	Replace population
			below_limit	Limit number of designs dominating those kept

Description

!!!!!!!!!! INCONSISTENCY in the manual. This says replacement types are the same between methods, however, this is not true in the input spec. below_limit vs. favor_feasible.

As of JEGA v2.0, all replacement types are common to both MOGA and SOGA. They include the `roulette_wheel`, `unique_roulette_wheel`, `elitist`, and `below_limit` selectors. In roulette_wheel replacement, each design is conceptually allotted a portion of a wheel proportional to its fitness relative to the fitnesses of the other Designs. Then, portions of the wheel are chosen at random and the design occupying those portions are duplicated into the next population. Those Designs allotted larger portions of the wheel are more likely to be selected (potentially many times). `unique_roulette_wheel` replacement is the same as `roulette_wheel` replacement, with the exception that a design may only be selected once. The `below_limit` selector attempts to keep all designs for which the negated fitness is below a certain limit. The values are negated to keep with the convention that higher fitness is better. The inputs to the `below_limit` selector are the limit as a real value, and a `shrinkage_percentage` as a real value. The `shrinkage_percentage` defines the minimum amount of selections that will take place if enough designs are available. It is interpreted as a percentage of the population size that must go on to the subsequent generation. To enforce this, `below_limit` makes all the selections it would make anyway and if that is not enough, it takes the remaining that it needs from the best of what is left (effectively raising its limit as far as it must to get the minimum number of selections). It continues until it has made enough selections. The `shrinkage_percentage` is designed to prevent extreme decreases in the population size at any given generation, and thus prevent a big loss of genetic diversity in a very short time. Without a shrinkage limit, a small group of "super" designs may appear and quickly cull the population down to a size on the order of the limiting value. In this case, all the diversity of the population is lost and it is expensive to re-diversify and spread the population. The `elitist` selector simply chooses the required number of designs taking the most fit. For example, if 100 selections are requested, then the top 100 designs as ranked by fitness will be selected and the remaining will be discarded.

The initialization, crossover, and mutation controls were all described above. There are no SOGA specific aspects to these controls. The `replacement_type` for a SOGA may be `roulette_wheel`, `unique_roulette_wheel`, `elitist`, or `favor_feasible`.

The `favor_feasible` replacement type first considers feasibility as a selection criteria. If that does not produce a "winner" then it moves on to considering fitness value. Because of this, any fitness assessor used with the `favor_feasible` selector must only account objectives in the creation of fitness. Therefore, there is such a fitness assessor and it's use is enforced when the \ favor_feasible selector is chosen. In that case, and if the output level is set high enough, a message will be presented indicating that the `weighted_sum_only` fitness assessor will be used. As of JEGA v2.0 and beyond, the fitness assessment operator must be specified with SOGA although the `merit_function` is currently the only one (note that the `weighted_sum_only` assessor exists but cannot be selected). The roulette wheel selectors no longer assume a fitness function. The `merit_function` fitness assessor uses an exterior penalty function formulation to penalize infeasible designs. The specification allows the input of a `constraint_penalty` which is the multiplier to use on the constraint violations.

elitist

- Keywords Area
- method
- moga
- replacement_type
- elitist

Use the best designs to form a new population

Specification

Alias: none
Argument(s): none

Description

The `elitist` (default) setting creates a new population using (a) the `replacement_size` best individuals from the current population, (b) and `population_size` - `replacement_size` individuals randomly selected from the newly generated individuals. It is possible in this case to lose a good solution from the newly generated individuals if it is not randomly selected for replacement; however, the default `new_solutions_generated` value is set such that the entire set of newly generated individuals will be selected for replacement.

roulette_wheel

- Keywords Area
- method
- moga
- replacement_type
- roulette_wheel

Replace population

Specification

Alias: none
Argument(s): none

Description

Replace the population with members selected to continue in the next generation. The pool of potential members is the current population and the current set of offspring. The `replacement_type` of `roulette_wheel` or `unique_roulette_wheel` may be used either with MOGA or SOGA problems however they are not recommended for use with MOGA. Given that the only two fitness assessors for MOGA are the `layer_rank` and `domination_count`, the recommended selector is the `below_limit` selector. The `below_limit` replacement will only keep designs that are dominated by fewer than a limiting number of other designs. The `replacement_type` of `favor_feasible` is specific to a SOGA. This replacement operator will always prefer a more feasible design to a less feasible one. Beyond that, it favors solutions based on an assigned fitness value which must have been installed by the weighted sum only fitness assessor (see the discussion below).

unique_roulette_wheel

- Keywords Area
- method
- moga
- replacement_type
- unique_roulette_wheel

Replace population

Specification

Alias: none
Argument(s): none

Description

Replace the population with members selected to continue in the next generation. The pool of potential members is the current population and the current set of offspring. The `replacement_type` of `roulette_wheel` or `unique_roulette_wheel` may be used either with MOGA or SOGA problems however they are not recommended for use with MOGA. Given that the only two fitness assessors for MOGA are the `layer_rank` and `domination_count`, the recommended selector is the `below_limit` selector. The `below_limit` replacement will only keep designs that are dominated by fewer than a limiting number of other designs. The `replacement_type` of `favor_feasible` is specific to a SOGA. This replacement operator will always prefer a more feasible design to a less feasible one. Beyond that, it favors solutions based on an assigned fitness value which must have been installed by the weighted sum only fitness assessor (see the discussion below).

below_limit

- Keywords Area
- method
- moga
- replacement_type
- below_limit

Limit number of designs dominating those kept

Specification

Alias: none
Argument(s): REAL

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		shrinkage_fraction	Decrease the population size by a percentage

Description

The `below_limit` replacement will only keep designs that are dominated by fewer than a limiting number of other designs.

shrinkage_fraction

- Keywords Area
- method
- moga
- replacement_type
- below_limit
- shrinkage_fraction

Decrease the population size by a percentage

Specification

Alias: shrinkage_percentage
Argument(s): REAL

Description

As of JEGA v2.0, all replacement types are common to both MOGA and SOGA. They include the `roulette_wheel`, `unique_roulette_wheel`, `elitist`, and `below_limit` selectors. In roulette_wheel replacement, each design is conceptually allotted a portion of a wheel proportional to its fitness relative to the fitnesses of the other Designs. Then, portions of the wheel are chosen at random and the design occupying those portions are duplicated into the next population. Those Designs allotted larger portions of the wheel are more likely to be selected (potentially many times). `unique_roulette_wheel` replacement is the same as `roulette_wheel` replacement, with the exception that a design may only be selected once. The `below_limit` selector attempts to keep all designs for which the negated fitness is below a certain limit. The values are negated to keep with the convention that higher fitness is better. The inputs to the `below_limit` selector are the limit as a real value, and a `shrinkage_percentage` as a real value. The `shrinkage_percentage` defines the minimum amount of selections that will take place if enough designs are available. It is interpreted as a percentage of the population size that must go on to the subsequent generation. To enforce this, `below_limit` makes all the selections it would make anyway and if that is not enough, it takes the remaining that it needs from the best of what is left (effectively raising its limit as far as it must to get the minimum number of selections). It continues until it has made enough selections. The `shrinkage_percentage` is designed to prevent extreme decreases in the population size at any given generation, and thus prevent a big loss of genetic diversity in a very short time. Without a shrinkage limit, a small group of "super" designs may appear and quickly cull the population down to a size on the order of the limiting value. In this case, all the diversity of the population is lost and it is expensive to re-diversify and spread the population. The `elitist` selector simply chooses the required number of designs taking the most fit. For example, if 100 selections are requested, then the top 100 designs as ranked by fitness will be selected and the remaining will be discarded.

niching_type

- Keywords Area
- method
- moga
- niching_type

Specify the type of niching pressure

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	radial	Set niching distance to percentage of non-dominated range

			distance	Enforce minimum Euclidean distance between designs
			max_designs	Limit number of solutions to remain in the population

Description

The purpose of niching is to encourage differentiation along the Pareto frontier and thus a more even and uniform sampling.

This is typically accomplished by discouraging clustering of design points in the performance space. In JEGA, the application of niche pressure occurs as a secondary selection operation. The nicher is given a chance to perform a pre-selection operation prior to the operation of the selection (replacement) operator, and is then called to perform niching on the set of designs that were selected by the selection operator.

The radial nicher takes information input from the user to compute a minimum allowable distance between designs in the performance space and acts as a secondary selection operator whereby it enforces this minimum distance. The distance nicher requires that solutions must be separated from other solutions by a minimum distance in each dimension (vs. Euclidean distance for the radial niching). After niching is complete, all designs in the population will be at least the minimum distance from one another in all directions.

The `radial` niche pressure applicator works by enforcing a minimum Euclidean distance between designs in the performance space at each generation. The algorithm proceeds by starting at the (or one of the) extreme designs along objective dimension 0 and marching through the population removing all designs that are too close to the current design. One exception to the rule is that the algorithm will never remove an extreme design which is defined as a design that is maximal or minimal in all but 1 objective dimension (for a classical 2 objective problem, the extreme designs are those at the tips of the non-dominated frontier). The `distance` nicher enforces a minimimum distance in each dimension.

The designs that are removed by the nicher are not discarded. They are buffered and re-inserted into the population during the next pre-selection operation. This way, the selector is still the only operator that discards designs and the algorithm will not waste time ”re-filling” gaps created by the nicher.

The `radial` nicher requires as input a vector of fractions with length equal to the number of objectives. The elements of the vector are interpreted as percentages of the non-dominated range for each objective defining a minimum distance to all other designs. All values should be in the range (0, 1). The minimum allowable distance between any two designs in the performance space is the Euclidian (simple square-root-sum-of-squares calculation) distance defined by these percentages. The `distance` nicher has a similar input vector requirement, only the distance is the minimum distance in each dimension.

The `max_designs` niche pressure applicator is designed to choose a limited number of solutions to remain in the population. That number is specified by `num_designs`. It does so in order to balance the tendency for populations to grow very large and thus consuming too many computer resources. It operates by ranking designs according to their fitness standing and a computed count of how many other designs are too close to them. Too close is a function of the supplied niche_vector, which specifies the minimum distance between any two points in the performance space along each dimension individually. Once the designs are all ranked, the top c\ num_-designs designs are kept in the population and the remaining ones are bufferred or discarded. Note that like other niching operators, this one will not discard an extreme design.

radial

- Keywords Area
- method

- moga
- niching_type
- radial

Set niching distance to percentage of non-dominated range

Specification

Alias: none
Argument(s): REALLIST

Description

The `radial` nicher requires as input a vector of fractions with length equal to the number of objectives. The elements of the vector are interpreted as percentages of the non-dominated range for each objective defining a minimum distance to all other designs. All values should be in the range (0, 1). The minimum allowable distance between any two designs in the performance space is the Euclidian (simple square-root-sum-of-squares calculation) distance defined by these percentages. The `distance` nicher has a similar input vector requirement, only the distance is the minimum distance in each dimension.

distance

- Keywords Area
- method
- moga
- niching_type
- distance

Enforce minimum Euclidean distance between designs

Specification

Alias: none
Argument(s): REALLIST

Description

Currently, the only niche pressure operators available are the `radial` nicher, the `distance` nicher, and the `max_designs` nicher. The `radial` niche pressure applicator works by enforcing a minimum Euclidean distance between designs in the performance space at each generation. The algorithm proceeds by starting at the (or one of the) extreme designs along objective dimension 0 and marching through the population removing all designs that are too close to the current design. One exception to the rule is that the algorithm will never remove an extreme design which is defined as a design that is maximal or minimal in all but 1 objective dimension (for a classical 2 objective problem, the extreme designs are those at the tips of the non-dominated frontier). The `distance` nicher enforces a minimimum distance in each dimension.

The designs that are removed by the nicher are not discarded. They are buffered and re-inserted into the population during the next pre-selection operation. This way, the selector is still the only operator that discards designs and the algorithm will not waste time ”re-filling” gaps created by the nicher.

max_designs

- Keywords Area
- method
- moga
- niching_type
- max_designs

Limit number of solutions to remain in the population

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_designs	Limit the number of solutions

Description

The `max_designs` niche pressure applicator is designed to choose a limited number of solutions to remain in the population. That number is specified by `num_designs`. It does so in order to balance the tendency for populations to grow very large and thus consuming too many computer resources. It operates by ranking designs according to their fitness standing and a computed count of how many other designs are too close to them. Too close is a function of the supplied niche_vector, which specifies the minimum distance between any two points in the performance space along each dimension individually. Once the designs are all ranked, the top c\ num_-designs designs are kept in the population and the remaining ones are bufferred or discarded. Note that like other niching operators, this one will not discard an extreme design.

num_designs

- Keywords Area
- method
- moga
- niching_type
- max_designs
- num_designs

Limit the number of solutions

Specification

Alias: none
Argument(s): INTEGER

Description

The `max_designs` niche pressure applicator is designed to choose a limited number of solutions to remain in the population. That number is specified by `num_designs`. It does so in order to balance the tendency for populations to grow very large and thus consuming too many computer resources. It operates by ranking designs according to their fitness standing and a computed count of how many other designs are too close to them. Too close is a function of the supplied niche_vector, which specifies the minimum distance between any two points in the performance space along each dimension individually. Once the designs are all ranked, the top c\ num_-designs designs are kept in the population and the remaining ones are bufferred or discarded. Note that like other niching operators, this one will not discard an extreme design.

convergence_type

- Keywords Area
- method
- moga
- convergence_type

Select the convergence type for JEGA methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		metric_tracker	Track changes in the non-dominated frontier
	Optional		percent_change	Define the convergence criterion for JEGA methods
	Optional		num_generations	Define the convergence criterion for JEGA methods

Description

The two JEGA methods use different convergence types, which are described on their respective pages.

All the convergence types are modified by the optional keywords `percent_change` and `num_generations`.

metric_tracker

- Keywords Area
- method

- moga
- convergence_type
- metric_tracker

Track changes in the non-dominated frontier

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): none

Description

The moga converger (`metric_tracker`) operates by tracking various changes in the non-dominated frontier from generation to generation. When the changes occurring over a user specified number of generations fall below a user specified threshold, the algorithm stops.

If `metric_tracker` is specified, then a `percent_change` and `num_generations` must be supplied as well. These are listed as optional keywords in the input spec.

Theory

The `metric_tracker` converger tracks 3 metrics specific to the non-dominated frontier from generation to generation. All 3 of these metrics are computed as percent changes between the generations. In order to compute these metrics, the converger stores a duplicate of the non-dominated frontier at each generation for comparison to the non-dominated frontier of the next generation.

The first metric is one that indicates how the expanse of the frontier is changing. The expanse along a given objective is defined by the range of values existing within the non-dominated set. The expansion metric is computed by tracking the extremes of the non-dominated frontier from one generation to the next. Any movement of the extreme values is noticed and the maximum percentage movement is computed as:

```
Em = max over j of abs((range(j, i) - range(j, i-1)) /
range(j, i-1)) j=1,nof
```

where Em is the max expansion metric, j is the objective function index, i is the current generation number, and nof is the total number of objectives. The range is the difference between the largest value along an objective and the smallest when considering only non-dominated designs.

The second metric monitors changes in the density of the non-dominated set. The density metric is computed as the number of non-dominated points divided by the hypervolume of the non-dominated region of space. Therefore, changes in the density can be caused by changes in the number of non-dominated points or by changes in size of the non-dominated space or both. The size of the non-dominated space is computed as:

```
  Vps(i) = product over j of range(j, i)  j=1,nof
```

where Vps(i) is the hypervolume of the non-dominated space at generation i and all other terms have the same meanings as above.

The density of the a given non-dominated space is then:

```
  Dps(i) = Pct(i) / Vps(i)
```

where Pct(i) is the number of points on the non-dominated frontier at generation i.

The percentage increase in density of the frontier is then calculated as

```
Cd = abs((Dps(i) - Dps(i-1)) / Dps(i-1))
```

where Cd is the change in density metric.

The final metric is one that monitors the "goodness" of the non-dominated frontier. This metric is computed by considering each design in the previous population and determining if it is dominated by any designs in the current population. All that are determined to be dominated are counted. The metric is the ratio of the number that are dominated to the total number that exist in the previous population.

As mentioned above, each of these metrics is a percentage. The tracker records the largest of these three at each generation. Once the recorded percentage is below the supplied percent change for the supplied number of generations consecutively, the algorithm is converged.

percent_change

- Keywords Area
- method
- moga
- convergence_type
- percent_change

Define the convergence criterion for JEGA methods

Specification

Alias: none

Argument(s): REAL

Description

The `percent_change` is the threshold beneath which convergence is attained whereby it is compared to the metric value computed.

num_generations

- Keywords Area
- method
- moga
- convergence_type
- num_generations

Define the convergence criterion for JEGA methods

Specification

Alias: none

Argument(s): INTEGER

Description

The `num_generations` is the number of generations over which the metric value should be tracked. Convergence will be attained if the recorded metric is below `percent_change` for `num_generations` consecutive generations.

postprocessor_type

- Keywords Area
- method
- moga
- postprocessor_type

Post process the final solution from `moga`

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		orthogonal_-distance	Get subset of Pareto front based on distance

Description

The purpose of this operation is to perform any needed data manipulations on the final solution deemed necessary. Currently the `orthogonal_distance` is the only one. It reduces the final solution set size such that a minimum distance in each direction exists between any two designs.

!!!!!!!!!!! Is this correct, the original did not match the input spec

orthogonal_distance

- Keywords Area
- method
- moga
- postprocessor_type
- orthogonal_distance

Get subset of Pareto front based on distance

Specification

Alias: none
Argument(s): REALLIST

Description

Note that MOGA and SOGA create additional output files during execution. "finaldata.dat" is a file that holds the final set of Pareto optimal solutions after any post-processing is complete. "discards.dat" holds solutions that were discarded from the population during the course of evolution. It can often be useful to plot objective function values from these files to visually see the Pareto front and ensure that finaldata.dat solutions dominate discards.dat solutions. The solutions are written to these output files in the format "Input1...InputN..Output1...OutputM". If MOGA is used in a hybrid optimization meta-iteration (which requires one optimal solution from each individual optimization method to be passed to the subsequent optimization method as its starting point), the solution in the Pareto set closest to the "utopia" point is given as the best solution. This solution is also reported in the Dakota output. This "best" solution in the Pareto set has minimum distance from the utopia point. The utopia point is defined as the point of extreme (best) values for each objective function. For example, if the Pareto front is bounded by (1,100) and (90,2), then (1,2) is the utopia point. There will be a point in the Pareto set that has minimum L2-norm distance to this point, for example (10,10) may be such a point. In SOGA, the solution that minimizes the single objective function is returned as the best solution. If moga is used in meta-iteration which may require passing multiple solutions to the next level (such as the `surrogate_based_global` or `hybrid` methods), the `orthogonal_distance` postprocessor type may be used to specify the distances between each solution value to winnow down the solutions in the full Pareto front to a subset which will be passed to the next iteration.

population_size

- Keywords Area
- method
- moga
- population_size

Set the initial population size in JEGA methods

Specification

Alias: none
Argument(s): INTEGER

Description

The number of designs in the initial population is specified by the `population_size`. Note that the `population_size` only sets the size of the initial population. The population size may vary in the JEGA methods according to the type of operators chosen for a particular optimization run.

log_file

- Keywords Area
- method
- moga
- log_file

Specify the name of a log file

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): STRING

Description

New as of JEGA v2.0 is the introduction of the `log_file` specification. JEGA now uses a logging library to output messages and status to the user. JEGA can be configured at build time to log to both the console window and a text file, one or the other, or neither. The `log_file` input is a string name of a file into which to log. If the build was configured without file logging in JEGA, this input is ignored. If file logging is enabled and no `log_file` is specified, the default file name of JEGAGlobal.log is used.

print_each_pop

- Keywords Area
- method
- moga
- print_each_pop

Print every population to a population file

Specification

Alias: none
Argument(s): none

Description

New to JEGA v2.0 is the introduction of the `print_each_pop` specification. It serves as a flag and if supplied, the population at each generation will be printed to a file named "population<GEN#>.dat" where <GEN#> is the number of the current generation.

initialization_type

- Keywords Area
- method
- moga
- initialization_type

Specify how to initialize the population

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	simple_random	Initilization option for JEGA methods
			unique_random	(Default) Initilization option for JEGA methods
			flat_file	Initilization option for JEGA methods

Description

The `initialization_type` defines how the initial population is created for the GA. There are three types:

1. `simple_random`
2. `unique_random` (default)
3. `flat_file`

Setting the size for the `flat_file` initializer has the effect of requiring a minimum number of designs to create. If this minimum number has not been created once the files are all read, the rest are created using the `unique_random` initializer and then the `simple_random` initializer if necessary.

simple_random

- Keywords Area
- method
- moga
- initialization_type
- simple_random

Initilization option for JEGA methods

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

`simple_random` creates initial solutions with random variable values according to a uniform random number distribution. It gives no consideration to any previously generated designs.

unique_random

- Keywords Area
- method
- moga
- initialization_type
- unique_random

(Default) Initilization option for JEGA methods

Specification

Alias: none
Argument(s): none

Description

`unique_random` is the same as `simple_random`, except that when a new solution is generated, it is checked against the rest of the solutions. If it duplicates any of them, it is rejected.

flat_file

- Keywords Area
- method
- moga
- initialization_type
- flat_file

Initilization option for JEGA methods

Specification

Alias: none
Argument(s): STRING

Description

`flat_file` allows the initial population to be read from a flat file. If `flat_file` is specified, a file name must be given.

Variables can be delimited in the flat file in any way you see fit with a few exceptions. The delimiter must be the same on any given line of input with the exception of leading and trailing whitespace. So a line could look like: 1.1, 2.2 ,3.3 for example but could not look like: 1.1, 2.2 3.3. The delimiter can vary from line to line within the file which can be useful if data from multiple sources is pasted into the same input file. The delimiter can be any string that does not contain any of the characters .+-dDeE or any of the digits 0-9. The input will be read until the end of the file. The algorithm will discard any configurations for which it was unable to retrieve at least the number of design variables. The objective and constraint entries are not required but if ALL are present, they will be recorded and the design will be tagged as evaluated so that evaluators may choose not to re-evaluate them.

Setting the size for this initializer has the effect of requiring a minimum number of designs to create. If this minimum number has not been created once the files are all read, the rest are created using the `unique_random` initializer and then the `simple_random` initializer if necessary.

crossover_type

- Keywords Area
- method
- moga
- crossover_type

Select a crossover type for JEGA methods

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** multi_point_binary	**Dakota Keyword Description** Use bit switching for crossover events
			multi_point_-parameterized_-binary	Use bit switching to crossover each design variable

			multi_point_real	Perform crossover in real valued genome
			shuffle_random	Perform crossover by choosing design variable(s)
	Optional		crossover_rate	Specify the probability of a crossover event

Description

There are many crossover types available. `multi_point_binary` crossover requires an integer number, N, of crossover points. This crossover type performs a bit switching crossover at N crossover points in the binary encoded genome of two designs. Thus, crossover may occur at any point along a solution chromosome (in the middle of a gene representing a design variable, for example). `multi_point_parameterized_binary` crossover is similar in that it performs a bit switching crossover routine at N crossover points. However, this crossover type performs crossover on each design variable individually. So the individual chromosomes are crossed at N locations. `multi_point_real` crossover performs a variable switching crossover routing at N crossover points in the real real valued genome of two designs. In this scheme, crossover only occurs between design variables (chromosomes). Note that the standard solution chromosome representation in the JEGA algorithm is real encoded and can handle integer or real design variables. For any crossover types that use a binary representation, real variables are converted to long integers by multiplying the real number by 10^6 and then truncating. Note that this assumes a precision of only six decimal places. Discrete variables are represented as integers (indices within a list of possible values) within the algorithm and thus require no special treatment by the binary operators.

The final crossover type is `shuffle_random`. This crossover type performs crossover by choosing design variables at random from a specified number of parents enough times that the requested number of children are produced. For example, consider the case of 3 parents producing 2 children. This operator would go through and for each design variable, select one of the parents as the donor for the child. So it creates a random shuffle of the parent design variable values. The relative numbers of children and parents are controllable to allow for as much mixing as desired. The more parents involved, the less likely that the children will wind up exact duplicates of the parents.

All crossover types take a `crossover_rate`. The crossover rate is used to calculate the number of crossover operations that take place. The number of crossovers is equal to the rate $*$ population_size.

multi_point_binary

- Keywords Area
- method
- moga
- crossover_type
- multi_point_binary

Use bit switching for crossover events

Specification

Alias: none

Argument(s): INTEGER

Description

There are many crossover types available. `multi_point_binary` crossover requires an integer number, N, of crossover points. This crossover type performs a bit switching crossover at N crossover points in the binary encoded genome of two designs. Thus, crossover may occur at any point along a solution chromosome (in the middle of a gene representing a design variable, for example). `multi_point_parameterized_binary` crossover is similar in that it performs a bit switching crossover routine at N crossover points. However, this crossover type performs crossover on each design variable individually. So the individual chromosomes are crossed at N locations. `multi_point_real` crossover performs a variable switching crossover routing at N crossover points in the real real valued genome of two designs. In this scheme, crossover only occurs between design variables (chromosomes). Note that the standard solution chromosome representation in the JEGA algorithm is real encoded and can handle integer or real design variables. For any crossover types that use a binary representation, real variables are converted to long integers by multiplying the real number by 10^6 and then truncating. Note that this assumes a precision of only six decimal places. Discrete variables are represented as integers (indices within a list of possible values) within the algorithm and thus require no special treatment by the binary operators.

multi_point_parameterized_binary

- Keywords Area
- method
- moga
- crossover_type
- multi_point_parameterized_binary

Use bit switching to crossover each design variable

Specification

Alias: none
Argument(s): INTEGER

Description

There are many crossover types available. `multi_point_binary` crossover requires an integer number, N, of crossover points. This crossover type performs a bit switching crossover at N crossover points in the binary encoded genome of two designs. Thus, crossover may occur at any point along a solution chromosome (in the middle of a gene representing a design variable, for example). `multi_point_parameterized_binary` crossover is similar in that it performs a bit switching crossover routine at N crossover points. However, this crossover type performs crossover on each design variable individually. So the individual chromosomes are crossed at N locations. `multi_point_real` crossover performs a variable switching crossover routing at N crossover points in the real real valued genome of two designs. In this scheme, crossover only occurs between design variables (chromosomes). Note that the standard solution chromosome representation in the JEGA algorithm is real encoded and can handle integer or real design variables. For any crossover types that use a binary representation, real variables are converted to long integers by multiplying the real number by 10^6 and then truncating. Note that this assumes a precision of only six decimal places. Discrete variables are represented as integers (indices within a list of possible values) within the algorithm and thus require no special treatment by the binary operators.

multi_point_real

- Keywords Area
- method
- moga
- crossover_type
- multi_point_real

Perform crossover in real valued genome

Specification

Alias: none
Argument(s): INTEGER

Description

There are many crossover types available. `multi_point_binary` crossover requires an integer number, N, of crossover points. This crossover type performs a bit switching crossover at N crossover points in the binary encoded genome of two designs. Thus, crossover may occur at any point along a solution chromosome (in the middle of a gene representing a design variable, for example). `multi_point_parameterized_binary` crossover is similar in that it performs a bit switching crossover routine at N crossover points. However, this crossover type performs crossover on each design variable individually. So the individual chromosomes are crossed at N locations. `multi_point_real` crossover performs a variable switching crossover routing at N crossover points in the real real valued genome of two designs. In this scheme, crossover only occurs between design variables (chromosomes). Note that the standard solution chromosome representation in the JEGA algorithm is real encoded and can handle integer or real design variables. For any crossover types that use a binary representation, real variables are converted to long integers by multiplying the real number by 10^6 and then truncating. Note that this assumes a precision of only six decimal places. Discrete variables are represented as integers (indices within a list of possible values) within the algorithm and thus require no special treatment by the binary operators.

shuffle_random

- Keywords Area
- method
- moga
- crossover_type
- shuffle_random

Perform crossover by choosing design variable(s)

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_parents	Number of parents in random shuffle crossover
	Optional		num_offspring	Number of offspring in random shuffle crossover

Description

The final crossover type is `shuffle_random`. This crossover type performs crossover by choosing design variables at random from a specified number of parents enough times that the requested number of children are produced. For example, consider the case of 3 parents producing 2 children. This operator would go through and for each design variable, select one of the parents as the donor for the child. So it creates a random shuffle of the parent design variable values. The relative numbers of children and parents are controllable to allow for as much mixing as desired. The more parents involved, the less likely that the children will wind up exact duplicates of the parents.

num_parents

- Keywords Area
- method
- moga
- crossover_type
- shuffle_random
- num_parents

Number of parents in random shuffle crossover

Specification

Alias: none
Argument(s): INTEGER

Description

Number of parents in random shuffle crossover

num_offspring

- Keywords Area
- method
- moga
- crossover_type

- shuffle_random
- num_offspring

Number of offspring in random shuffle crossover

Specification

Alias: none
Argument(s): INTEGER

Description

Number of offspring in random shuffle crossover

crossover_rate

- Keywords Area
- method
- moga
- crossover_type
- crossover_rate

Specify the probability of a crossover event

Specification

Alias: none
Argument(s): REAL

Description

The `crossover_type` controls what approach is employed for combining parent genetic information to create offspring, and the `crossover_rate` specifies the probability of a crossover operation being performed to generate a new offspring. The SCOLIB EA method supports three forms of crossover, `two_point`, `blend`, and `uniform`, which generate a new individual through combinations of two parent individuals. Two-point crossover divides each parent into three regions, where offspring are created from the combination of the middle region from one parent and the end regions from the other parent. Since the SCOLIB EA does not utilize bit representations of variable values, the crossover points only occur on coordinate boundaries, never within the bits of a particular coordinate. Uniform crossover creates offspring through random combination of coordinates from the two parents. Blend crossover generates a new individual randomly along the multidimensional vector connecting the two parents.

mutation_type

- Keywords Area
- method
- moga
- mutation_type

Select a mutation type for JEGA methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	bit_random	Mutate by flipping a random bit
			replace_uniform	Use uniformly distributed value over range of parameter
			offset_normal	Set mutation offset to use a normal distribution
			offset_cauchy	Use a Cauchy distribution for the mutation offset
			offset_uniform	Set mutation offset to use a uniform distribution
	Optional		mutation_rate	Set probability of a mutation

Description

There are five mutation types allowed. `replace_uniform` introduces random variation by first randomly choosing a design variable of a randomly selected design and reassigning it to a random valid value for that variable. No consideration of the current value is given when determining the new value. All mutation types have a `mutation_rate`. The number of mutations for the replace_uniform mutator is the product of the mutation_rate and the population_size.

The `bit_random` mutator introduces random variation by first converting a randomly chosen variable of a randomly chosen design into a binary string. It then flips a randomly chosen bit in the string from a 1 to a 0 or visa versa. In this mutation scheme, the resulting value has more probability of being similar to the original value. The number of mutations performed is the product of the mutation_rate, the number of design variables, and the population size.

The offset mutators all act by adding an ”offset” random amount to a variable value. The random amount has a mean of zero in all cases. The `offset_normal` mutator introduces random variation by adding a Gaussian random amount to a variable value. The random amount has a standard deviation dependent on the `mutation_scale`. The `mutation_scale` is a fraction in the range [0, 1] and is meant to help control the amount of

variation that takes place when a variable is mutated. `mutation_scale` is multiplied by the range of the variable being mutated to serve as standard deviation. `offset_cauchy` is similar to `offset_normal`, except that a Cauchy random variable is added to the variable being mutated. The `mutation_scale` also defines the standard deviation for this mutator. Finally, `offset_uniform` adds a uniform random amount to the variable value. For the `offset_uniform` mutator, the `mutation_scale` is interpreted as a fraction of the total range of the variable. The range of possible deviation amounts is +/- 1/2 ∗ (`mutation_scale` ∗ variable range). The number of mutations for all offset mutators is defined as the product of `mutation_rate` and `population_size`.

bit_random

- Keywords Area
- method
- moga
- mutation_type
- bit_random

Mutate by flipping a random bit

Specification

Alias: none
Argument(s): none

Description

The `bit_random` mutator introduces random variation by first converting a randomly chosen variable of a randomly chosen design into a binary string. It then flips a randomly chosen bit in the string from a 1 to a 0 or visa versa. In this mutation scheme, the resulting value has more probability of being similar to the original value. The number of mutations performed is the product of the mutation_rate, the number of design variables, and the population size.

replace_uniform

- Keywords Area
- method
- moga
- mutation_type
- replace_uniform

Use uniformly distributed value over range of parameter

Specification

Alias: none
Argument(s): none

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

offset_normal

- Keywords Area
- method
- moga
- mutation_type
- offset_normal

Set mutation offset to use a normal distribution

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		mutation_scale	Scales mutation across range of parameter

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from

the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

mutation_scale

- Keywords Area
- method
- moga
- mutation_type
- offset_normal
- mutation_scale

Scales mutation across range of parameter

Specification

Alias: none
Argument(s): REAL

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

Note that mutation_scale and mutation_range are not recognized by Dakota as valid keywords unless mutation_type has been specified and the type is an "offset" mutations.

offset_cauchy

- Keywords Area
- method
- moga
- mutation_type
- offset_cauchy

Use a Cauchy distribution for the mutation offset

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		mutation_scale	Scales mutation across range of parameter

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

mutation_scale

- Keywords Area
- method

- moga
- mutation_type
- offset_cauchy
- mutation_scale

Scales mutation across range of parameter

Specification

Alias: none
Argument(s): REAL

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

Note that `mutation_scale` and `mutation_range` are not recognized by Dakota as valid keywords unless `mutation_type` has been specified and the type is an "offset" mutations.

offset_uniform

- Keywords Area
- method
- moga
- mutation_type
- offset_uniform

Set mutation offset to use a uniform distribution

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		mutation_scale	Scales mutation across range of parameter

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

mutation_scale

- Keywords Area
- method
- moga
- mutation_type
- offset_uniform
- mutation_scale

Scales mutation across range of parameter

Specification

Alias: none
Argument(s): REAL

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given

coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

Note that `mutation_scale` and `mutation_range` are not recognized by Dakota as valid keywords unless `mutation_type` has been specified and the type is an "offset" mutations.

mutation_rate

- Keywords Area
- method
- moga
- mutation_type
- mutation_rate

Set probability of a mutation

Specification

Alias: none
Argument(s): REAL

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These

offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

seed

- Keywords Area
- method
- moga
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

linear_inequality_constraint_matrix

- Keywords Area
- method
- moga
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- moga
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- moga
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- moga
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- moga
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\mathrm{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \mathrm{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- moga
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- moga
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of `0.` for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- moga
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- moga
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$
$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- moga
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.38 soga

- Keywords Area
- method
- soga

Single-objective Genetic Algorithm (a.k.a Evolutionary Algorithm)

Topics

This keyword is related to the topics:

- package_jega
- global_optimization_methods

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		fitness_type	Select the fitness type for JEGA methods
	Optional		replacement_type	Select a replacement type for JEGA methods

	Optional	convergence_type	Select the convergence type for JEGA methods
	Optional	population_size	Set the initial population size in JEGA methods
	Optional	log_file	Specify the name of a log file
	Optional	print_each_pop	Print every population to a population file
	Optional	initialization_type	Specify how to initialize the population
	Optional	crossover_type	Select a crossover type for JEGA methods
	Optional	mutation_type	Select a mutation type for JEGA methods
	Optional	seed	Specify the seed of the random number generator
	Optional	linear_inequality_-constraint_matrix	Define coefficients of the linear inequality constraints
	Optional	linear_inequality_-lower_bounds	Define lower bounds for the linear inequality constraint
	Optional	linear_inequality_-upper_bounds	Define upper bounds for the linear inequality constraint
	Optional	linear_inequality_-scale_types	Specify how each linear inequality constraint is scaled
	Optional	linear_inequality_-scales	Define the characteristic values to scale linear inequalities

	Optional	linear_equality_-constraint_matrix	Define coefficients of the linear equalities
	Optional	linear_equality_-targets	Define target values for the linear equality constraints
	Optional	linear_equality_-scale_types	Specify how each linear equality constraint is scaled
	Optional	linear_equality_-scales	Define the characteristic values to scale linear equalities
	Optional	model_pointer	Specify which model block will be used by a method

Description

`soga` stands for Single-objective Genetic Algorithm, which is a global optimization method that supports general constraints and a mixture of real and discrete variables. `soga` is part of the JEGA library.

Constraints `soga` can utilize linear constraints.

Configuration

The genetic algorithm configurations are:

1. fitness
2. replacement
3. convergence
4. initialization
5. crossover
6. mutation
7. population size

The pool of potential members is the current population and the current set of offspring. Choice of fitness assessors is strongly related to the type of replacement algorithm being used and can have a profound effect on the solutions selected for the next generation.

Stopping Criteria

The `soga` method respects the `max_iterations` and `max_function_evaluations` method independent controls to provide integer limits for the maximum number of generations and function evaluations, respectively.

The algorithm also stops when convergence is reached. This involves repeated assessment of the algorithm's progress in solving the problem, until some criterion is met.

Outputs The `soga` method respects the `output` method independent control to vary the amount of information presented to the user during execution.

The final results are written to the Dakota tabular output. Additional information is also available - see the `log_file` and `print_each_pop` keywords.

Theory

The basic steps of the `soga` algorithm are as follows:

1. Initialize the population
2. Evaluate the population (calculate the values of the objective function and constraints for each population member)
3. Loop until converged, or stopping criteria reached
   (a) Perform crossover
   (b) Perform mutation
   (c) Evaluate the new population
   (d) Assess the fitness of each member in the population
   (e) Replace the population with members selected to continue in the next generation
   (f) Test for convergence

See Also

These keywords may also be of interest:

- moga

fitness_type

- Keywords Area
- method
- soga
- fitness_type

Select the fitness type for JEGA methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		merit_function	Balance goals of reducing objective function and satisfying constraints

	Optional	constraint_penalty	Multiplier for the penalty function

Description

The two JEGA methods use different fitness types, which are described on their respective pages.

merit_function

- Keywords Area
- method
- soga
- fitness_type
- merit_function

Balance goals of reducing objective function and satisfying constraints

Specification

Alias: none
Argument(s): none

Description

A `merit_function` is a function in constrained optimization that attempts to provide joint progress toward reducing the objective function and satisfying the constraints.

constraint_penalty

- Keywords Area
- method
- soga
- fitness_type
- constraint_penalty

Multiplier for the penalty function

Specification

Alias: none
Argument(s): REAL

Description

- `constraint_penalty`: the initial multiplier for the penalty function, must be nonnegative

Most SCOLIB optimizers treat constraints with a simple penalty scheme that adds `constraint_penalty` times the sum of squares of the constraint violations to the objective function. Specific exceptions to this method for handling constraint violations are noted below. (The default value of `constraint_penalty` is 1000.0, except for methods that dynamically adapt their constraint penalty, for which the default value is 1.0.)

replacement_type

- Keywords Area
- method
- soga
- replacement_type

Select a replacement type for JEGA methods

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** elitist	**Dakota Keyword Description** Use the best designs to form a new population
			favor_feasible	Prioritize feasible designs
			roulette_wheel	Replace population
			unique_roulette_-wheel	Replace population

Description

!!!!!!!!!! INCONSISTENCY in the manual. This says replacement types are the same between methods, however, this is not true in the input spec. below_limit vs. favor_feasible.

As of JEGA v2.0, all replacement types are common to both MOGA and SOGA. They include the `roulette-_wheel`, `unique_roulette_wheel`, `elitist`, and `below_limit` selectors. In roulette_wheel replacement, each design is conceptually allotted a portion of a wheel proportional to its fitness relative to the fitnesses of the other Designs. Then, portions of the wheel are chosen at random and the design occupying those portions are duplicated into the next population. Those Designs allotted larger portions of the wheel are more likely to be

selected (potentially many times). unique_roulette_wheel replacement is the same as roulette_wheel replacement, with the exception that a design may only be selected once. The below_limit selector attempts to keep all designs for which the negated fitness is below a certain limit. The values are negated to keep with the convention that higher fitness is better. The inputs to the below_limit selector are the limit as a real value, and a shrinkage_percentage as a real value. The shrinkage_percentage defines the minimum amount of selections that will take place if enough designs are available. It is interpreted as a percentage of the population size that must go on to the subsequent generation. To enforce this, below_limit makes all the selections it would make anyway and if that is not enough, it takes the remaining that it needs from the best of what is left (effectively raising its limit as far as it must to get the minimum number of selections). It continues until it has made enough selections. The shrinkage_percentage is designed to prevent extreme decreases in the population size at any given generation, and thus prevent a big loss of genetic diversity in a very short time. Without a shrinkage limit, a small group of "super" designs may appear and quickly cull the population down to a size on the order of the limiting value. In this case, all the diversity of the population is lost and it is expensive to re-diversify and spread the population. The elitist selector simply chooses the required number of designs taking the most fit. For example, if 100 selections are requested, then the top 100 designs as ranked by fitness will be selected and the remaining will be discarded.

The initialization, crossover, and mutation controls were all described above. There are no SOGA specific aspects to these controls. The replacement_type for a SOGA may be roulette_wheel, unique_roulette_wheel, elitist, or favor_feasible.

The favor_feasible replacement type first considers feasibility as a selection criteria. If that does not produce a "winner" then it moves on to considering fitness value. Because of this, any fitness assessor used with the favor_feasible selector must only account objectives in the creation of fitness. Therefore, there is such a fitness assessor and it's use is enforced when the \ favor_feasible selector is chosen. In that case, and if the output level is set high enough, a message will be presented indicating that the weighted_sum_only fitness assessor will be used. As of JEGA v2.0 and beyond, the fitness assessment operator must be specified with SOGA although the merit_function is currently the only one (note that the weighted_sum_only assessor exists but cannot be selected). The roulette wheel selectors no longer assume a fitness function. The merit_function fitness assessor uses an exterior penalty function formulation to penalize infeasible designs. The specification allows the input of a constraint_penalty which is the multiplier to use on the constraint violations.

elitist

- Keywords Area
- method
- soga
- replacement_type
- elitist

Use the best designs to form a new population

Specification

Alias: none

Argument(s): none

Description

The `elitist` (default) setting creates a new population using (a) the `replacement_size` best individuals from the current population, (b) and `population_size` - `replacement_size` individuals randomly selected from the newly generated individuals. It is possible in this case to lose a good solution from the newly generated individuals if it is not randomly selected for replacement; however, the default `new_solutions_generated` value is set such that the entire set of newly generated individuals will be selected for replacement.

favor_feasible

- Keywords Area
- method
- soga
- replacement_type
- favor_feasible

Prioritize feasible designs

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): none

Description

This replacement operator will always prefer a more feasible design to a less feasible one. Beyond that, it favors solutions based on an assigned fitness value which must have been installed by the weighted sum only fitness assessor (see the discussion below).

roulette_wheel

- Keywords Area
- method
- soga
- replacement_type
- roulette_wheel

Replace population

Specification

Alias: none
Argument(s): none

Description

Replace the population with members selected to continue in the next generation. The pool of potential members is the current population and the current set of offspring. The `replacement_type` of `roulette_wheel` or `unique_roulette_wheel` may be used either with MOGA or SOGA problems however they are not recommended for use with MOGA. Given that the only two fitness assessors for MOGA are the `layer_rank` and `domination_count`, the recommended selector is the `below_limit` selector. The `below_limit` replacement will only keep designs that are dominated by fewer than a limiting number of other designs. The `replacement_type` of `favor_feasible` is specific to a SOGA. This replacement operator will always prefer a more feasible design to a less feasible one. Beyond that, it favors solutions based on an assigned fitness value which must have been installed by the weighted sum only fitness assessor (see the discussion below).

unique_roulette_wheel

- Keywords Area
- method
- soga
- replacement_type
- unique_roulette_wheel

Replace population

Specification

Alias: none
Argument(s): none

Description

Replace the population with members selected to continue in the next generation. The pool of potential members is the current population and the current set of offspring. The `replacement_type` of `roulette_wheel` or `unique_roulette_wheel` may be used either with MOGA or SOGA problems however they are not recommended for use with MOGA. Given that the only two fitness assessors for MOGA are the `layer_rank` and `domination_count`, the recommended selector is the `below_limit` selector. The `below_limit` replacement will only keep designs that are dominated by fewer than a limiting number of other designs. The `replacement_type` of `favor_feasible` is specific to a SOGA. This replacement operator will always prefer a more feasible design to a less feasible one. Beyond that, it favors solutions based on an assigned fitness value which must have been installed by the weighted sum only fitness assessor (see the discussion below).

convergence_type

- Keywords Area
- method
- soga
- convergence_type

Select the convergence type for JEGA methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	best_fitness_tracker	Tracks the best fitness of the population
			average_fitness_-tracker	Tracks the average fitness of the population

Description

The two JEGA methods use different convergence types, which are described on their respective pages.

All the convergence types are modified by the optional keywords `percent_change` and `num_generations`.

best_fitness_tracker

- Keywords Area
- method
- soga
- convergence_type
- best_fitness_tracker

Tracks the best fitness of the population

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Optional	percent_change	Define the convergence criterion for JEGA methods
	Optional	num_generations	Define the convergence criterion for JEGA methods

Description

The `best_fitness_tracker` tracks the best fitness in the population. Convergence occurs after `num_generations` has passed and there has been less than `percent_change` in the best fitness value. The percent change can be as low as 0% in which case there must be no change at all over the number of generations.

See Also

These keywords may also be of interest:

- average_fitness_tracker

percent_change

- Keywords Area
- method
- soga
- convergence_type
- best_fitness_tracker
- percent_change

Define the convergence criterion for JEGA methods

Specification

Alias: none
Argument(s): REAL

Description

The `percent_change` is the threshold beneath which convergence is attained whereby it is compared to the metric value computed.

num_generations

- Keywords Area
- method
- soga
- convergence_type
- best_fitness_tracker
- num_generations

Define the convergence criterion for JEGA methods

Specification

Alias: none
Argument(s): INTEGER

Description

The `num_generations` is the number of generations over which the metric value should be tracked. Convergence will be attained if the recorded metric is below `percent_change` for `num_generations` consecutive generations.

average_fitness_tracker

- Keywords Area
- method
- soga
- convergence_type
- average_fitness_tracker

Tracks the average fitness of the population

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		percent_change	Define the convergence criterion for JEGA methods

	Optional	num_generations	Define the convergence criterion for JEGA methods

Description

The `convergence_type` called `average_fitness_tracker` keeps track of the average fitness in a population. If this average fitness does not change more than `percent_change` over some number of generations, `num_generations`, then the solution is reported as converged and the algorithm terminates.

See Also

These keywords may also be of interest:

- best_fitness_tracker

percent_change

- Keywords Area
- method
- soga
- convergence_type
- average_fitness_tracker
- percent_change

Define the convergence criterion for JEGA methods

Specification

Alias: none
Argument(s): REAL

Description

The `percent_change` is the threshold beneath which convergence is attained whereby it is compared to the metric value computed.

num_generations

- Keywords Area
- method
- soga
- convergence_type
- average_fitness_tracker
- num_generations

Define the convergence criterion for JEGA methods

Specification

Alias: none
Argument(s): INTEGER

Description

The `num_generations` is the number of generations over which the metric value should be tracked. Convergence will be attained if the recorded metric is below `percent_change` for `num_generations` consecutive generations.

population_size

- Keywords Area
- method
- soga
- population_size

Set the initial population size in JEGA methods

Specification

Alias: none
Argument(s): INTEGER

Description

The number of designs in the initial population is specified by the `population_size`. Note that the `population_size` only sets the size of the initial population. The population size may vary in the JEGA methods according to the type of operators chosen for a particular optimization run.

log_file

- Keywords Area
- method
- soga
- log_file

Specify the name of a log file

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): STRING

Description

New as of JEGA v2.0 is the introduction of the `log_file` specification. JEGA now uses a logging library to output messages and status to the user. JEGA can be configured at build time to log to both the console window and a text file, one or the other, or neither. The `log_file` input is a string name of a file into which to log. If the build was configured without file logging in JEGA, this input is ignored. If file logging is enabled and no `log_file` is specified, the default file name of JEGAGlobal.log is used.

print_each_pop

- Keywords Area
- method
- soga
- print_each_pop

Print every population to a population file

Specification

Alias: none
Argument(s): none

Description

New to JEGA v2.0 is the introduction of the `print_each_pop` specification. It serves as a flag and if supplied, the population at each generation will be printed to a file named "population<GEN#>.dat" where <GEN#> is the number of the current generation.

initialization_type

- Keywords Area
- method
- soga
- initialization_type

Specify how to initialize the population

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	simple_random	Initilization option for JEGA methods
			unique_random	(Default) Initilization option for JEGA methods
			flat_file	Initilization option for JEGA methods

Description

The `initialization_type` defines how the initial population is created for the GA. There are three types:

1. `simple_random`
2. `unique_random` (default)
3. `flat_file`

Setting the size for the `flat_file` initializer has the effect of requiring a minimum number of designs to create. If this minimum number has not been created once the files are all read, the rest are created using the `unique_random` initializer and then the `simple_random` initializer if necessary.

simple_random

- Keywords Area
- method
- soga
- initialization_type
- simple_random

Initilization option for JEGA methods

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none

Argument(s): none

Description

`simple_random` creates initial solutions with random variable values according to a uniform random number distribution. It gives no consideration to any previously generated designs.

unique_random

- Keywords Area
- method
- soga
- initialization_type
- unique_random

(Default) Initilization option for JEGA methods

Specification

Alias: none
Argument(s): none

Description

`unique_random` is the same as `simple_random`, except that when a new solution is generated, it is checked against the rest of the solutions. If it duplicates any of them, it is rejected.

flat_file

- Keywords Area
- method
- soga
- initialization_type
- flat_file

Initilization option for JEGA methods

Specification

Alias: none
Argument(s): STRING

Description

`flat_file` allows the initial population to be read from a flat file. If `flat_file` is specified, a file name must be given.

Variables can be delimited in the flat file in any way you see fit with a few exceptions. The delimiter must be the same on any given line of input with the exception of leading and trailing whitespace. So a line could look like: 1.1, 2.2 ,3.3 for example but could not look like: 1.1, 2.2 3.3. The delimiter can vary from line to line within the file which can be useful if data from multiple sources is pasted into the same input file. The delimiter can be any string that does not contain any of the characters .+-dDeE or any of the digits 0-9. The input will be read until the end of the file. The algorithm will discard any configurations for which it was unable to retrieve at least the number of design variables. The objective and constraint entries are not required but if ALL are present, they will be recorded and the design will be tagged as evaluated so that evaluators may choose not to re-evaluate them.

Setting the size for this initializer has the effect of requiring a minimum number of designs to create. If this minimum number has not been created once the files are all read, the rest are created using the `unique_random` initializer and then the `simple_random` initializer if necessary.

crossover_type

- Keywords Area
- method
- soga
- crossover_type

Select a crossover type for JEGA methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	multi_point_binary	Use bit switching for crossover events
			multi_point_-parameterized_-binary	Use bit switching to crossover each design variable
			multi_point_real	Perform crossover in real valued genome
			shuffle_random	Perform crossover by choosing design variable(s)

	Optional	crossover_rate	Specify the probability of a crossover event

Description

There are many crossover types available. `multi_point_binary` crossover requires an integer number, N, of crossover points. This crossover type performs a bit switching crossover at N crossover points in the binary encoded genome of two designs. Thus, crossover may occur at any point along a solution chromosome (in the middle of a gene representing a design variable, for example). `multi_point_parameterized_binary` crossover is similar in that it performs a bit switching crossover routine at N crossover points. However, this crossover type performs crossover on each design variable individually. So the individual chromosomes are crossed at N locations. `multi_point_real` crossover performs a variable switching crossover routing at N crossover points in the real real valued genome of two designs. In this scheme, crossover only occurs between design variables (chromosomes). Note that the standard solution chromosome representation in the JEGA algorithm is real encoded and can handle integer or real design variables. For any crossover types that use a binary representation, real variables are converted to long integers by multiplying the real number by 10^6 and then truncating. Note that this assumes a precision of only six decimal places. Discrete variables are represented as integers (indices within a list of possible values) within the algorithm and thus require no special treatment by the binary operators.

The final crossover type is `shuffle_random`. This crossover type performs crossover by choosing design variables at random from a specified number of parents enough times that the requested number of children are produced. For example, consider the case of 3 parents producing 2 children. This operator would go through and for each design variable, select one of the parents as the donor for the child. So it creates a random shuffle of the parent design variable values. The relative numbers of children and parents are controllable to allow for as much mixing as desired. The more parents involved, the less likely that the children will wind up exact duplicates of the parents.

All crossover types take a `crossover_rate`. The crossover rate is used to calculate the number of crossover operations that take place. The number of crossovers is equal to the rate $*$ population_size.

multi_point_binary

- Keywords Area
- method
- soga
- crossover_type
- multi_point_binary

Use bit switching for crossover events

Specification

Alias: none

Argument(s): INTEGER

Description

There are many crossover types available. `multi_point_binary` crossover requires an integer number, N, of crossover points. This crossover type performs a bit switching crossover at N crossover points in the binary encoded genome of two designs. Thus, crossover may occur at any point along a solution chromosome (in the middle of a gene representing a design variable, for example). `multi_point_parameterized_binary` crossover is similar in that it performs a bit switching crossover routine at N crossover points. However, this crossover type performs crossover on each design variable individually. So the individual chromosomes are crossed at N locations. `multi_point_real` crossover performs a variable switching crossover routing at N crossover points in the real real valued genome of two designs. In this scheme, crossover only occurs between design variables (chromosomes). Note that the standard solution chromosome representation in the JEGA algorithm is real encoded and can handle integer or real design variables. For any crossover types that use a binary representation, real variables are converted to long integers by multiplying the real number by 10^6 and then truncating. Note that this assumes a precision of only six decimal places. Discrete variables are represented as integers (indices within a list of possible values) within the algorithm and thus require no special treatment by the binary operators.

multi_point_parameterized_binary

- Keywords Area
- method
- soga
- crossover_type
- multi_point_parameterized_binary

Use bit switching to crossover each design variable

Specification

Alias: none
Argument(s): INTEGER

Description

There are many crossover types available. `multi_point_binary` crossover requires an integer number, N, of crossover points. This crossover type performs a bit switching crossover at N crossover points in the binary encoded genome of two designs. Thus, crossover may occur at any point along a solution chromosome (in the middle of a gene representing a design variable, for example). `multi_point_parameterized_binary` crossover is similar in that it performs a bit switching crossover routine at N crossover points. However, this crossover type performs crossover on each design variable individually. So the individual chromosomes are crossed at N locations. `multi_point_real` crossover performs a variable switching crossover routing at N crossover points in the real real valued genome of two designs. In this scheme, crossover only occurs between design variables (chromosomes). Note that the standard solution chromosome representation in the JEGA algorithm is real encoded and can handle integer or real design variables. For any crossover types that use a binary representation, real variables are converted to long integers by multiplying the real number by 10^6 and then truncating. Note that this assumes a precision of only six decimal places. Discrete variables are represented as integers (indices within a list of possible values) within the algorithm and thus require no special treatment by the binary operators.

multi_point_real

- Keywords Area
- method
- soga
- crossover_type
- multi_point_real

Perform crossover in real valued genome

Specification

Alias: none
Argument(s): INTEGER

Description

There are many crossover types available. `multi_point_binary` crossover requires an integer number, N, of crossover points. This crossover type performs a bit switching crossover at N crossover points in the binary encoded genome of two designs. Thus, crossover may occur at any point along a solution chromosome (in the middle of a gene representing a design variable, for example). `multi_point_parameterized_binary` crossover is similar in that it performs a bit switching crossover routine at N crossover points. However, this crossover type performs crossover on each design variable individually. So the individual chromosomes are crossed at N locations. `multi_point_real` crossover performs a variable switching crossover routing at N crossover points in the real real valued genome of two designs. In this scheme, crossover only occurs between design variables (chromosomes). Note that the standard solution chromosome representation in the JEGA algorithm is real encoded and can handle integer or real design variables. For any crossover types that use a binary representation, real variables are converted to long integers by multiplying the real number by 10^6 and then truncating. Note that this assumes a precision of only six decimal places. Discrete variables are represented as integers (indices within a list of possible values) within the algorithm and thus require no special treatment by the binary operators.

shuffle_random

- Keywords Area
- method
- soga
- crossover_type
- shuffle_random

Perform crossover by choosing design variable(s)

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_parents	Number of parents in random shuffle crossover
	Optional		num_offspring	Number of offspring in random shuffle crossover

Description

The final crossover type is `shuffle_random`. This crossover type performs crossover by choosing design variables at random from a specified number of parents enough times that the requested number of children are produced. For example, consider the case of 3 parents producing 2 children. This operator would go through and for each design variable, select one of the parents as the donor for the child. So it creates a random shuffle of the parent design variable values. The relative numbers of children and parents are controllable to allow for as much mixing as desired. The more parents involved, the less likely that the children will wind up exact duplicates of the parents.

num_parents

- Keywords Area
- method
- soga
- crossover_type
- shuffle_random
- num_parents

Number of parents in random shuffle crossover

Specification

Alias: none
Argument(s): INTEGER

Description

Number of parents in random shuffle crossover

num_offspring

- Keywords Area
- method
- soga
- crossover_type

- shuffle_random
- num_offspring

Number of offspring in random shuffle crossover

Specification

Alias: none
Argument(s): INTEGER

Description

Number of offspring in random shuffle crossover

crossover_rate

- Keywords Area
- method
- soga
- crossover_type
- crossover_rate

Specify the probability of a crossover event

Specification

Alias: none
Argument(s): REAL

Description

The `crossover_type` controls what approach is employed for combining parent genetic information to create offspring, and the `crossover_rate` specifies the probability of a crossover operation being performed to generate a new offspring. The SCOLIB EA method supports three forms of crossover, `two_point`, `blend`, and `uniform`, which generate a new individual through combinations of two parent individuals. Two-point crossover divides each parent into three regions, where offspring are created from the combination of the middle region from one parent and the end regions from the other parent. Since the SCOLIB EA does not utilize bit representations of variable values, the crossover points only occur on coordinate boundaries, never within the bits of a particular coordinate. Uniform crossover creates offspring through random combination of coordinates from the two parents. Blend crossover generates a new individual randomly along the multidimensional vector connecting the two parents.

mutation_type

- Keywords Area
- method
- soga
- mutation_type

Select a mutation type for JEGA methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	bit_random	Mutate by flipping a random bit
			replace_uniform	Use uniformly distributed value over range of parameter
			offset_normal	Set mutation offset to use a normal distribution
			offset_cauchy	Use a Cauchy distribution for the mutation offset
			offset_uniform	Set mutation offset to use a uniform distribution
	Optional		mutation_rate	Set probability of a mutation

Description

There are five mutation types allowed. `replace_uniform` introduces random variation by first randomly choosing a design variable of a randomly selected design and reassigning it to a random valid value for that variable. No consideration of the current value is given when determining the new value. All mutation types have a `mutation_rate`. The number of mutations for the replace_uniform mutator is the product of the mutation_rate and the population_size.

The `bit_random` mutator introduces random variation by first converting a randomly chosen variable of a randomly chosen design into a binary string. It then flips a randomly chosen bit in the string from a 1 to a 0 or visa versa. In this mutation scheme, the resulting value has more probability of being similar to the original value. The number of mutations performed is the product of the mutation_rate, the number of design variables, and the population size.

The offset mutators all act by adding an ”offset” random amount to a variable value. The random amount has a mean of zero in all cases. The `offset_normal` mutator introduces random variation by adding a Gaussian random amount to a variable value. The random amount has a standard deviation dependent on the `mutation_scale`. The `mutation_scale` is a fraction in the range [0, 1] and is meant to help control the amount of

variation that takes place when a variable is mutated. `mutation_scale` is multiplied by the range of the variable being mutated to serve as standard deviation. `offset_cauchy` is similar to `offset_normal`, except that a Cauchy random variable is added to the variable being mutated. The `mutation_scale` also defines the standard deviation for this mutator. Finally, `offset_uniform` adds a uniform random amount to the variable value. For the `offset_uniform` mutator, the `mutation_scale` is interpreted as a fraction of the total range of the variable. The range of possible deviation amounts is +/- 1/2 * (`mutation_scale` * variable range). The number of mutations for all offset mutators is defined as the product of `mutation_rate` and `population_size`.

bit_random

- Keywords Area
- method
- soga
- mutation_type
- bit_random

Mutate by flipping a random bit

Specification

Alias: none
Argument(s): none

Description

The `bit_random` mutator introduces random variation by first converting a randomly chosen variable of a randomly chosen design into a binary string. It then flips a randomly chosen bit in the string from a 1 to a 0 or visa versa. In this mutation scheme, the resulting value has more probability of being similar to the original value. The number of mutations performed is the product of the mutation_rate, the number of design variables, and the population size.

replace_uniform

- Keywords Area
- method
- soga
- mutation_type
- replace_uniform

Use uniformly distributed value over range of parameter

Specification

Alias: none
Argument(s): none

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

offset_normal

- Keywords Area
- method
- soga
- mutation_type
- offset_normal

Set mutation offset to use a normal distribution

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		mutation_scale	Scales mutation across range of parameter

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from

the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

mutation_scale

- Keywords Area
- method
- soga
- mutation_type
- offset_normal
- mutation_scale

Scales mutation across range of parameter

Specification

Alias: none
Argument(s): REAL

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

Note that mutation_scale and mutation_range are not recognized by Dakota as valid keywords unless mutation_type has been specified and the type is an "offset" mutations.

offset_cauchy

- Keywords Area
- method
- soga
- mutation_type
- offset_cauchy

Use a Cauchy distribution for the mutation offset

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		mutation_scale	Scales mutation across range of parameter

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

mutation_scale

- Keywords Area
- method

- soga
- mutation_type
- offset_cauchy
- mutation_scale

Scales mutation across range of parameter

Specification

Alias: none
Argument(s): REAL

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

Note that `mutation_scale` and `mutation_range` are not recognized by Dakota as valid keywords unless `mutation_type` has been specified and the type is an "offset" mutations.

offset_uniform

- Keywords Area
- method
- soga
- mutation_type
- offset_uniform

Set mutation offset to use a uniform distribution

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		mutation_scale	Scales mutation across range of parameter

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

mutation_scale

- Keywords Area
- method
- soga
- mutation_type
- offset_uniform
- mutation_scale

Scales mutation across range of parameter

Specification

Alias: none
Argument(s): REAL

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given

coordinate value with a random variable (a ”replace” mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are ”offset” mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

Note that mutation_scale and mutation_range are not recognized by Dakota as valid keywords unless mutation_type has been specified and the type is an ”offset” mutations.

mutation_rate

- Keywords Area
- method
- soga
- mutation_type
- mutation_rate

Set probability of a mutation

Specification

Alias: none
Argument(s): REAL

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an ”offset” mutation), but also by replacing a given coordinate value with a random variable (a ”replace” mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are ”offset” mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These

offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

seed

- Keywords Area
- method
- soga
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

linear_inequality_constraint_matrix

- Keywords Area
- method
- soga
- linear_inequality_constraint_matrix

Define coefficients of the linear inequality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the constraint matrix A provides coefficients for the variables in the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where the bounds are optionally specified by `linear_inequality_lower_bounds`, and `linear_inequality_upper_bounds`. The bounds, if not specified, will default to -infinity, and 0, respectively, resulting in one-sided inequalities of the form

$$Ax \leq 0.0$$

.

linear_inequality_lower_bounds

- Keywords Area
- method
- soga
- linear_inequality_lower_bounds

Define lower bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_upper_bounds

- Keywords Area
- method
- soga
- linear_inequality_upper_bounds

Define upper bounds for the linear inequality constraint

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the inequality case, the lower a_l and upper a_u bounds provide constraint limits for the two-sided formulation:

$$a_l \leq Ax \leq a_u$$

Where A is the constrain matrix of variable coefficients.

As with nonlinear inequality constraints (see objective_functions), the default linear inequality constraint bounds are selected so that one-sided inequalities of the form

$$Ax \leq 0.0$$

result when there are no user bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity.

This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`).

linear_inequality_scale_types

- Keywords Area
- method
- soga
- linear_inequality_scale_types

Specify how each linear inequality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_inequality_scale_types` provide strings specifying the scaling type for each linear inequality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_inequality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_inequality_scales

- Keywords Area
- method
- soga
- linear_inequality_scales

Define the characteristic values to scale linear inequalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_inequality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

Behavior depends on the choice of `linear_inequality_scale_type` :

- `scale_type` - behavior of `linear_inequality_scales`
- `'none'` - ignored
- `'value'` - required
- `'auto'` - optional

If a single real value is specified it will apply to all components of the constraint.

Scaling for linear constraints is applied *after* any continuous variable scaling. For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_constraint_matrix

- Keywords Area
- method
- soga
- linear_equality_constraint_matrix

Define coefficients of the linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the constraint matrix A provides coefficients for the variables on the left hand side of:

$$Ax = a_t$$

linear_equality_targets

- Keywords Area
- method
- soga
- linear_equality_targets

Define target values for the linear equality constraints

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

In the equality case, the targets a_t provide the equality constraint right hand sides:

$$Ax = a_t$$

.

If this is not specified, the defaults for the equality constraint targets enforce a value of $0.$ for each constraint: $Ax = 0.0$

linear_equality_scale_types

- Keywords Area
- method
- soga
- linear_equality_scale_types

Specify how each linear equality constraint is scaled

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): STRINGLIST

Description

`linear_equality_scale_types` provide strings specifying the scaling type for each linear equality constraint, in methods that support scaling.

An entry may be selected for each constraint. The options are:

- `'none'` - no scaling
- `'value'` - characteristic value if this is chosen, then `linear_equality_scales` must be specified
- `'auto'` - automatic scaling If a single string is specified it will apply to all constraints.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear equality constraints

$$a_L \leq A_i x \leq a_U$$

$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$

$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$

$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

linear_equality_scales

- Keywords Area
- method
- soga
- linear_equality_scales

Define the characteristic values to scale linear equalities

Topics

This keyword is related to the topics:

- linear_constraints

Specification

Alias: none
Argument(s): REALLIST

Description

Each entry in `linear_equality_scales` may be a user-specified, nonzero characteristic value to be used in scaling each constraint.

See the scaling page for details on how to use this keyword.

Scaling for linear constraints is applied *after* any continuous variable scaling.

For example, for variable scaling on continuous design variables x:

$$\tilde{x}^j = \frac{x^j - x_O^j}{x_M^j}$$

we have the following system for linear inequality constraints

$$a_L \leq A_i x \leq a_U$$
$$a_L \leq A_i \left(\text{diag}(x_M)\tilde{x} + x_O\right) \leq a_U$$
$$a_L - A_i x_O \leq A_i \text{diag}(x_M)\tilde{x} \leq a_U - A_i x_O$$
$$\tilde{a}_L \leq \tilde{A}_i \tilde{x} \leq \tilde{a}_U$$

and user-specified or automatically computed scaling multipliers are appplied to this final transformed system, which accounts for continuous design variable scaling. When automatic scaling is in use for linear constraints they are linearly scaled by a computed characteristic value, but not affinely to [0,1].

model_pointer

- Keywords Area
- method
- soga
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.39 coliny_pattern_search

- Keywords Area
- method
- coliny_pattern_search

Pattern search, derivative free optimization method

Topics

This keyword is related to the topics:

- package_scolib
- package_coliny
- global_optimization_methods

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		constant_penalty	Use a simple weighted penalty to manage feasibility

	Optional	no_expansion	Don't allow expansion of the search pattern
	Optional	expand_after_-success	Set the factor by which a search pattern can be expanded
	Optional	pattern_basis	Pattern basis selection
	Optional	stochastic	Generate trial points in random order
	Optional	total_pattern_size	Total number of points in search pattern
	Optional	exploratory_moves	Exploratory moves selection
	Optional	synchronization	Select how Dakota schedules function evaluations
	Optional	contraction_factor	Amount by which step length is rescaled
	Optional	constraint_penalty	Multiplier for the penalty function
	Optional	initial_delta	Initial step size for non-gradient based optimizers
	Optional	threshold_delta	Stopping criteria based on step length or pattern size
	Optional	solution_target	Stopping criteria based on objective function value
	Optional	seed	Specify the seed of the random number generator
	Optional	show_misc_options	Show algorithm parameters not exposed in Dakota input
	Optional	misc_options	Set method options not available through Dakota spec
	Optional	model_pointer	Specify which model block will be used by a method

Description

Pattern search techniques are nongradient-based optimization methods which use a set of offsets from the current iterate to locate improved points in the design space.

See the page package_scolib for important information regarding all SCOLIB methods

`coliny_pattern_search` supports concurrency up to the size of the search pattern

Traditional pattern search methods search with a fixed pattern of search directions to try to find improvements to the current iterate. The SCOLIB pattern search methods generalize this simple algorithmic strategy to enable control of how the search pattern is adapted, as well as how each search pattern is evaluated. The `stochastic` and `synchronization` specifications denote how the the trial points are evaluated. The `stochastic` specification indicates that the trial points are considered in a random order. For parallel pattern search, `synchronization` dictates whether the evaluations are scheduled using a `blocking` scheduler or a `nonblocking` scheduler. In the `blocking` case, all points in the pattern are evaluated (in parallel), and if the best of these trial points is an improving point, then it becomes the next iterate. These runs are reproducible, assuming use of the same seed in the `stochastic` case. In the `nonblocking` case, all points in the pattern may not be evaluated, since the first improving point found becomes the next iterate. Since the algorithm steps will be subject to parallel timing variabilities, these runs will not generally be repeatable. The `synchronization` specification has similar connotations for sequential pattern search. If `blocking` is specified, then each sequential iteration terminates after all trial points have been considered, and if `nonblocking` is specified, then each sequential iteration terminates after the first improving trial point is evaluated. In this release, both `blocking` and `nonblocking` specifications result in blocking behavior (except in the case where `exporatory_moves` below is set to `adaptive_pattern`). Nonblocking behavior will be re-enabled after some underlying technical issues have been resolved.

The particular form of the search pattern is controlled by the `pattern_basis` specification. If `pattern_basis` is `coordinate` basis, then the pattern search uses a plus and minus offset in each coordinate direction, for a total of *2n* function evaluations in the pattern. This case is depicted in Figure 5.3 for three coordinate dimensions.

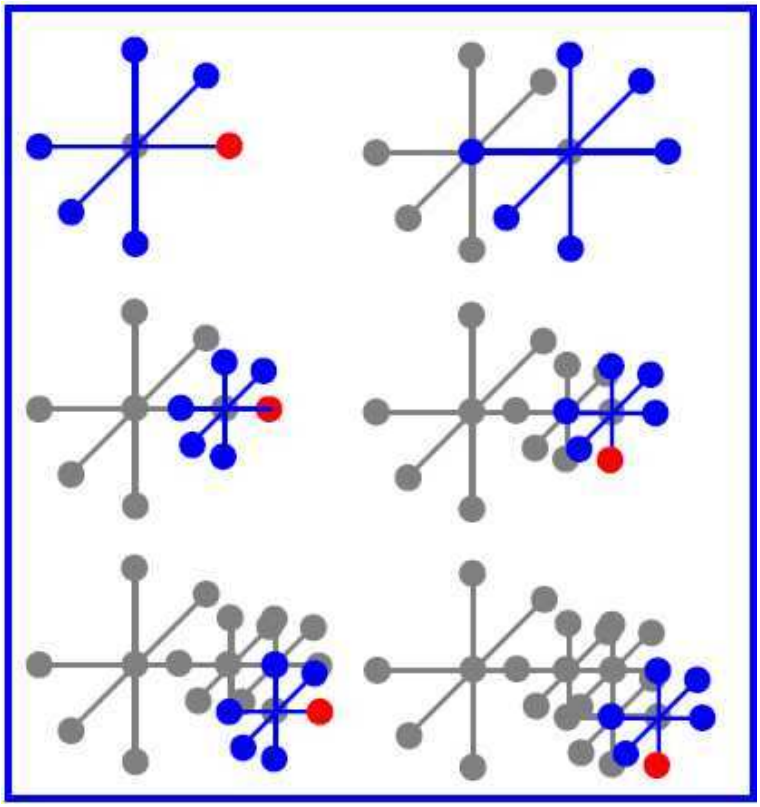

Figure 5.1: Depiction of coordinate pattern search algorithm

If `pattern_basis` is `simplex`, then pattern search uses a minimal positive basis simplex for the parameter space, for a total of *n+1* function evaluations in the pattern. Note that the `simplex` pattern basis can be used for unbounded problems only. The `total_pattern_size` specification can be used to augment the basic `coordinate` and `simplex` patterns with additional function evaluations, and is particularly useful for parallel load balancing. For example, if some function evaluations in the pattern are dropped due to duplication or bound constraint interaction, then the `total_pattern_size` specification instructs the algorithm to generate new offsets to bring the total number of evaluations up to this consistent total.

The `exploratory_moves` specification controls how the search pattern is adapted. (The search pattern can be adapted after an improving trial point is found, or after all trial points in a search pattern have been found to be unimproving points.) The following exploratory moves selections are supported by SCOLIB:

- The `basic_pattern` case is the simple pattern search approach, which uses the same pattern in each iteration.

- The `multi_step` case examines each trial step in the pattern in turn. If a successful step is found, the pattern search continues examining trial steps about this new point. In this manner, the effects of multiple successful steps are cumulative within a single iteration. This option does not support any parallelism and will result in a serial pattern search.

- The `adaptive_pattern` case invokes a pattern search technique that adaptively rescales the different search directions to maximize the number of redundant function evaluations. See [[45] "Hart et al., 2001"]

for details of this method. In preliminary experiments, this method had more robust performance than the standard basic_pattern case in serial tests. This option supports a limited degree of parallelism. After successful iterations (where the step length is not contracted), a parallel search will be performed. After unsuccessful iterations (where the step length is contracted), only a single evaluation is performed.

The initial_delta and threshold_delta specifications provide the initial offset size and the threshold size at which to terminate the algorithm. For any dimension that has both upper and lower bounds, this step length will be internally rescaled to provide search steps of length initial_delta $*$ range $*$ 0.1. This rescaling does not occur for other dimensions, so search steps in those directions have length initial_delta. Note that the factor of 0.1 in the rescaling could result in an undesirably small initial step. This can be offset by providing a large initial_delta.

In general, pattern search methods can expand and contract their step lengths. SCOLIB pattern search methods contract the step length by the value contraction_factor, and they expand the step length by the value (1/contraction_factor). The expand_after_success control specifies how many successful objective function improvements must occur with a specific step length prior to expansion of the step length, whereas the no_expansion flag instructs the algorithm to forgo pattern expansion altogether.

Finally, constraint infeasibility can be managed in a somewhat more sophisticated manner than the simple weighted penalty function. If the constant_penalty specification is used, then the simple weighted penalty scheme described above is used. Otherwise, the constraint penalty is adapted to the value constraint_penalty/L, where L is the the smallest step length used so far.

See Also

These keywords may also be of interest:

- coliny_beta
- coliny_direct
- coliny_cobyla
- coliny_ea
- coliny_solis_wets

constant_penalty

- Keywords Area
- method
- coliny_pattern_search
- constant_penalty

Use a simple weighted penalty to manage feasibility

Specification

Alias: none

Argument(s): none

Description

Finally, constraint infeasibility can be managed in a somewhat more sophisticated manner than the simple weighted penalty function. If the `constant_penalty` specification is used, then the simple weighted penalty scheme described above is used. Otherwise, the constraint penalty is adapted to the value `constraint_penalty/L`, where L is the the smallest step length used so far.

no_expansion

- Keywords Area
- method
- coliny_pattern_search
- no_expansion

Don't allow expansion of the search pattern

Specification

Alias: none
Argument(s): none

Description

In general, pattern search methods can expand and contract their step lengths. SCOLIB pattern search methods contract the step length by the value `contraction_factor`, and they expand the step length by the value (1/contraction_factor). The `expand_after_success` control specifies how many successful objective function improvements must occur with a specific step length prior to expansion of the step length, whereas the `no_expansion` flag instructs the algorithm to forgo pattern expansion altogether.

expand_after_success

- Keywords Area
- method
- coliny_pattern_search
- expand_after_success

Set the factor by which a search pattern can be expanded

Specification

Alias: none
Argument(s): INTEGER

Description

In general, pattern search methods can expand and contract their step lengths. SCOLIB pattern search methods contract the step length by the value `contraction_factor`, and they expand the step length by the value (1/contraction_factor). The `expand_after_success` control specifies how many successful objective function improvements must occur with a specific step length prior to expansion of the step length, whereas the `no_expansion` flag instructs the algorithm to forgo pattern expansion altogether.

pattern_basis

- Keywords Area
- method
- coliny_pattern_search
- pattern_basis

Pattern basis selection

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** coordinate	**Dakota Keyword Description** Use coordinate directions as search pattern
			simplex	Use a minimal simplex for the search pattern

Description

The particular form of the search pattern is controlled by the `pattern_basis` specification. If `pattern_basis` is `coordinate` basis, then the pattern search uses a plus and minus offset in each coordinate direction, for a total of $2n$ function evaluations in the pattern. This case is depicted in Figure 5.3 for three coordinate dimensions.

coordinate

- Keywords Area
- method
- coliny_pattern_search
- pattern_basis
- coordinate

Use coordinate directions as search pattern

Specification

Alias: none
Argument(s): none

Description

The particular form of the search pattern is controlled by the `pattern_basis` specification. If `pattern_basis` is `coordinate` basis, then the pattern search uses a plus and minus offset in each coordinate direction, for a total of *2n* function evaluations in the pattern. This case is depicted in Figure 5.3 for three coordinate dimensions.

simplex

- Keywords Area
- method
- coliny_pattern_search
- pattern_basis
- simplex

Use a minimal simplex for the search pattern

Specification

Alias: none
Argument(s): none

Description

If `pattern_basis` is `simplex`, then pattern search uses a minimal positive basis simplex for the parameter space, for a total of *n+1* function evaluations in the pattern. Note that the `simplex` pattern basis can be used for unbounded problems only. The `total_pattern_size` specification can be used to augment the basic `coordinate` and `simplex` patterns with additional function evaluations, and is particularly useful for parallel load balancing. For example, if some function evaluations in the pattern are dropped due to duplication or bound constraint interaction, then the `total_pattern_size` specification instructs the algorithm to generate new offsets to bring the total number of evaluations up to this consistent total.

stochastic

- Keywords Area
- method
- coliny_pattern_search
- stochastic

Generate trial points in random order

Specification

Alias: none
Argument(s): none

Description

Traditional pattern search methods search with a fixed pattern of search directions to try to find improvements to the current iterate. The SCOLIB pattern search methods generalize this simple algorithmic strategy to enable control of how the search pattern is adapted, as well as how each search pattern is evaluated. The `stochastic` and `synchronization` specifications denote how the the trial points are evaluated. The `stochastic` specification indicates that the trial points are considered in a random order. For parallel pattern search, `synchronization` dictates whether the evaluations are scheduled using a `blocking` scheduler or a `nonblocking` scheduler (i.e.,

total_pattern_size

- Keywords Area
- method
- coliny_pattern_search
- total_pattern_size

Total number of points in search pattern

Specification

Alias: none
Argument(s): INTEGER

Description

If `pattern_basis` is `simplex`, then pattern search uses a minimal positive basis simplex for the parameter space, for a total of *n+1* function evaluations in the pattern. Note that the `simplex` pattern basis can be used for unbounded problems only. The `total_pattern_size` specification can be used to augment the basic `coordinate` and `simplex` patterns with additional function evaluations, and is particularly useful for parallel load balancing. For example, if some function evaluations in the pattern are dropped due to duplication or bound constraint interaction, then the `total_pattern_size` specification instructs the algorithm to generate new offsets to bring the total number of evaluations up to this consistent total.

exploratory_moves

- Keywords Area
- method
- coliny_pattern_search
- exploratory_moves

Exploratory moves selection

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	Description of Group **Group 1**	Dakota Keyword multi_step	Dakota Keyword Description Examine trial step around successful new point
			adaptive_pattern	Adaptively rescale search directions
			basic_pattern	Use the same search pattern every iteration

Description

The `exploratory_moves` specification controls how the search pattern is adapted. (The search pattern can be adapted after an improving trial point is found, or after all trial points in a search pattern have been found to be unimproving points.) The following exploratory moves selections are supported by SCOLIB:

multi_step

- Keywords Area
- method
- coliny_pattern_search
- exploratory_moves
- multi_step

Examine trial step around successful new point

Specification

Alias: none
Argument(s): none

Description

The `multi_step` case examines each trial step in the pattern in turn. If a successful step is found, the pattern search continues examining trial steps about this new point. In this manner, the effects of multiple successful steps are cumulative within a single iteration. This option does not support any parallelism and will result in a serial pattern

adaptive_pattern

- Keywords Area
- method

- coliny_pattern_search
- exploratory_moves
- adaptive_pattern

Adaptively rescale search directions

Specification

Alias: none
Argument(s): none

Description

The `adaptive_pattern` case invokes a pattern search technique that adaptively rescales the different search directions to maximize the number of redundant function evaluations. See [[45] "Hart et al., 2001"] for details of this method. In preliminary experiments, this method had more robust performance than the standard `basic_pattern` case in serial tests. This option supports a limited degree of parallelism. After successful iterations (where the step length is not contracted), a parallel search will be performed. After unsuccessful iterations (where the step length is contracted), only a single evaluation is performed.

basic_pattern

- Keywords Area
- method
- coliny_pattern_search
- exploratory_moves
- basic_pattern

Use the same search pattern every iteration

Specification

Alias: none
Argument(s): none

Description

The `basic_pattern` case is the simple pattern search approach, which uses the same pattern in each iteration.

synchronization

- Keywords Area
- method
- coliny_pattern_search
- synchronization

Select how Dakota schedules function evaluations

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	blocking	The `blocking` option causes Dakota to run the algorithm in synchronous mode
			nonblocking	The `nonblocking` option causes Dakota to run the algorithm in asynchronous mode

Description

The `synchronization` specification can be used to specify the use of either `blocking` or `nonblocking` schedulers.

blocking

- Keywords Area
- method
- coliny_pattern_search
- synchronization
- blocking

The `blocking` option causes Dakota to run the algorithm in synchronous mode

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none

Argument(s): none

Description

The `blocking` option causes Dakota to run the algorithm in synchronous mode.

nonblocking

- Keywords Area
- method
- coliny_pattern_search
- synchronization
- nonblocking

The `nonblocking` option causes Dakota to run the algorithm in asynchronous mode

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The `nonblocking` option causes Dakota to run the algorithm in asynchronous mode.
The `nonblocking` option is not available when Dakota is used in message-passing mode.

contraction_factor

- Keywords Area
- method
- coliny_pattern_search
- contraction_factor

Amount by which step length is rescaled

Specification

Alias: none
Argument(s): REAL

Description

For pattern search methods, `contraction_factor` specifies the amount by which step length is rescaled after unsuccesful iterates, must be strictly between 0 and 1.

For methods that can expand the step length, the expansion is 1/ `contraction_factor`

constraint_penalty

- Keywords Area
- method
- coliny_pattern_search
- constraint_penalty

Multiplier for the penalty function

Specification

Alias: none

Argument(s): REAL

Description

- `constraint_penalty`: the initial multiplier for the penalty function, must be nonnegative

Most SCOLIB optimizers treat constraints with a simple penalty scheme that adds `constraint_penalty` times the sum of squares of the constraint violations to the objective function. Specific exceptions to this method for handling constraint violations are noted below. (The default value of `constraint_penalty` is 1000.0, except for methods that dynamically adapt their constraint penalty, for which the default value is 1.0.)

initial_delta

- Keywords Area
- method
- coliny_pattern_search
- initial_delta

Initial step size for non-gradient based optimizers

Specification

Alias: none

Argument(s): REAL

Description

If `initial_delta` is supplied by the user, it will be applied in an absolute sense in all coordinate directions. APPS documentation advocates choosing `initial_delta` to be the approximate distance from the initial point to the solution. If this is unknown, it is advisable to err on the side of choosing an `initial_delta` that is too large or to not specify it. In the latter case, APPS will take a full step to the boundary in each direction. Relative application of `initial_delta` is not available unless the user scales the problem accordingly.

threshold_delta

- Keywords Area
- method
- coliny_pattern_search
- threshold_delta

Stopping criteria based on step length or pattern size

Specification

Alias: none
Argument(s): REAL

Description

`threshold_delta` is the step length or pattern size used to determine convergence.

solution_target

- Keywords Area
- method
- coliny_pattern_search
- solution_target

Stopping criteria based on objective function value

Topics

This keyword is related to the topics:

- problem

Specification

Alias: solution_accuracy
Argument(s): REAL

Description

`solution_target` is a termination criterion. The algorithm will terminate when the function value falls below solution_target.

seed

- Keywords Area
- method
- coliny_pattern_search
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

show_misc_options

- Keywords Area
- method
- coliny_pattern_search
- show_misc_options

Show algorithm parameters not exposed in Dakota input

Specification

Alias: none
Argument(s): none

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

misc_options

- Keywords Area
- method
- coliny_pattern_search
- misc_options

Set method options not available through Dakota spec

Specification

Alias: none
Argument(s): STRINGLIST

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

model_pointer

- Keywords Area
- method
- coliny_pattern_search
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.40 coliny_solis_wets

- Keywords Area
- method
- coliny_solis_wets

Simple greedy local search method

Topics

This keyword is related to the topics:

- package_scolib
- package_coliny

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		contract_after_-failure	The number of unsuccessful cycles prior to contraction.
	Optional		no_expansion	Don't allow expansion of the search pattern
	Optional		expand_after_-success	Set the factor by which a search pattern can be expanded
	Optional		constant_penalty	Use a simple weighted penalty to manage feasibility

	Optional	contraction_factor	Amount by which step length is rescaled
	Optional	constraint_penalty	Multiplier for the penalty function
	Optional	initial_delta	Initial step size for non-gradient based optimizers
	Optional	threshold_delta	Stopping criteria based on step length or pattern size
	Optional	solution_target	Stopping criteria based on objective function value
	Optional	seed	Specify the seed of the random number generator
	Optional	show_misc_options	Show algorithm parameters not exposed in Dakota input
	Optional	misc_options	Set method options not available through Dakota spec
	Optional	model_pointer	Specify which model block will be used by a method

Description

The Solis-Wets method is a simple greedy local search heuristic for continuous parameter spaces. Solis-Wets generates trial points using a multivariate normal distribution, and unsuccessful trial points are reflected about the current point to find a descent direction.

See the page package_scolib for important information regarding all SCOLIB methods

`coliny_solis_wets` is inherently serial, no concurrency is used.

These specifications have the same meaning as corresponding specifications for coliny_pattern_search. Please see that page for specification details.

In particular, `coliny_solis_wets` supports dynamic rescaling of the step length, and dynamic rescaling of the constraint penalty. The only new specification is `contract_after_failure`, which specifies the number of unsuccessful cycles which must occur with a specific delta prior to contraction of the delta.

See Also

These keywords may also be of interest:

- coliny_beta
- coliny_direct

- coliny_pattern_search
- coliny_cobyla
- coliny_ea

contract_after_failure

- Keywords Area
- method
- coliny_solis_wets
- contract_after_failure

The number of unsuccessful cycles prior to contraction.

Specification

Alias: none
Argument(s): INTEGER

Description

In particular, `coliny_solis_wets` supports dynamic rescaling of the step length, and dynamic rescaling of the constraint penalty. The only new specification is `contract_after_failure`, which specifies the number of unsuccessful cycles which must occur with a specific delta prior to contraction of the delta.

no_expansion

- Keywords Area
- method
- coliny_solis_wets
- no_expansion

Don't allow expansion of the search pattern

Specification

Alias: none
Argument(s): none

Description

In general, pattern search methods can expand and contract their step lengths. SCOLIB pattern search methods contract the step length by the value `contraction_factor`, and they expand the step length by the value (1/contraction_factor). The `expand_after_success` control specifies how many successful objective function improvements must occur with a specific step length prior to expansion of the step length, whereas the `no_expansion` flag instructs the algorithm to forgo pattern expansion altogether.

expand_after_success

- Keywords Area
- method
- coliny_solis_wets
- expand_after_success

Set the factor by which a search pattern can be expanded

Specification

Alias: none
Argument(s): INTEGER

Description

In general, pattern search methods can expand and contract their step lengths. SCOLIB pattern search methods contract the step length by the value `contraction_factor`, and they expand the step length by the value (1/contraction_factor). The `expand_after_success` control specifies how many successful objective function improvements must occur with a specific step length prior to expansion of the step length, whereas the `no_expansion` flag instructs the algorithm to forgo pattern expansion altogether.

constant_penalty

- Keywords Area
- method
- coliny_solis_wets
- constant_penalty

Use a simple weighted penalty to manage feasibility

Specification

Alias: none
Argument(s): none

Description

Finally, constraint infeasibility can be managed in a somewhat more sophisticated manner than the simple weighted penalty function. If the `constant_penalty` specification is used, then the simple weighted penalty scheme described above is used. Otherwise, the constraint penalty is adapted to the value `constraint_penalty/L`, where L is the the smallest step length used so far.

contraction_factor

- Keywords Area
- method
- coliny_solis_wets
- contraction_factor

Amount by which step length is rescaled

Specification

Alias: none
Argument(s): REAL

Description

For pattern search methods, `contraction_factor` specifies the amount by which step length is rescaled after unsuccesful iterates, must be strictly between 0 and 1.

For methods that can expand the step length, the expansion is 1/ `contraction_factor`

constraint_penalty

- Keywords Area
- method
- coliny_solis_wets
- constraint_penalty

Multiplier for the penalty function

Specification

Alias: none
Argument(s): REAL

Description

- `constraint_penalty`: the initial multiplier for the penalty function, must be nonnegative

Most SCOLIB optimizers treat constraints with a simple penalty scheme that adds `constraint_penalty` times the sum of squares of the constraint violations to the objective function. Specific exceptions to this method for handling constraint violations are noted below. (The default value of `constraint_penalty` is 1000.0, except for methods that dynamically adapt their constraint penalty, for which the default value is 1.0.)

initial_delta

- Keywords Area
- method
- coliny_solis_wets
- initial_delta

Initial step size for non-gradient based optimizers

Specification

Alias: none
Argument(s): REAL

Description

If `initial_delta` is supplied by the user, it will be applied in an absolute sense in all coordinate directions. APPS documentation advocates choosing `initial_delta` to be the approximate distance from the initial point to the solution. If this is unknown, it is advisable to err on the side of choosing an `initial_delta` that is too large or to not specify it. In the latter case, APPS will take a full step to the boundary in each direction. Relative application of `initial_delta` is not available unless the user scales the problem accordingly.

threshold_delta

- Keywords Area
- method
- coliny_solis_wets
- threshold_delta

Stopping criteria based on step length or pattern size

Specification

Alias: none
Argument(s): REAL

Description

`threshold_delta` is the step length or pattern size used to determine convergence.

solution_target

- Keywords Area
- method
- coliny_solis_wets
- solution_target

Stopping criteria based on objective function value

Topics

This keyword is related to the topics:

- problem

Specification

Alias: solution_accuracy
Argument(s): REAL

Description

`solution_target` is a termination criterion. The algorithm will terminate when the function value falls below solution_target.

seed

- Keywords Area
- method
- coliny_solis_wets
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

show_misc_options

- Keywords Area
- method
- coliny_solis_wets
- show_misc_options

Show algorithm parameters not exposed in Dakota input

Specification

Alias: none
Argument(s): none

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

misc_options

- Keywords Area
- method
- coliny_solis_wets
- misc_options

Set method options not available through Dakota spec

Specification

Alias: none
Argument(s): STRINGLIST

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

model_pointer

- Keywords Area
- method
- coliny_solis_wets
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.41 coliny_cobyla

- Keywords Area
- method
- coliny_cobyla

Constrained Optimization BY Linear Approximations (COBYLA)

Topics

This keyword is related to the topics:

- package_scolib
- package_coliny
- local_optimization_methods
- constrained

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_delta	Initial step size for non-gradient based optimizers

	Optional	threshold_delta	Stopping criteria based on step length or pattern size
	Optional	solution_target	Stopping criteria based on objective function value
	Optional	seed	Specify the seed of the random number generator
	Optional	show_misc_options	Show algorithm parameters not exposed in Dakota input
	Optional	misc_options	Set method options not available through Dakota spec
	Optional	model_pointer	Specify which model block will be used by a method

Description

The Constrained Optimization BY Linear Approximations (COBYLA) algorithm is an extension to the Nelder--Mead simplex algorithm for handling general linear/nonlinear constraints and is invoked using the `coliny_cobyla` group specification. The COBYLA algorithm employs linear approximations to the objective and constraint functions, the approximations being formed by linear interpolation at N+1 points in the space of the variables. We regard these interpolation points as vertices of a simplex. The step length parameter controls the size of the simplex and it is reduced automatically from `initial_delta` to `threshold_delta`. One advantage that COBYLA has over many of its competitors is that it treats each constraint individually when calculating a change to the variables, instead of lumping the constraints together into a single penalty function.

See the page package_scolib for important information regarding all SCOLIB methods

`coliny_cobyla` is inherently serial.

Stopping Critieria

DIRECT can be terminated with:

- max_function_evaluations
- `solution_target`

COBYLA currently only supports termination based on the `max_function_evaluations` and `solution_target` specifications.

See Also

These keywords may also be of interest:

- coliny_beta
- coliny_direct

- coliny_pattern_search
- coliny_ea
- coliny_solis_wets

initial_delta

- Keywords Area
- method
- coliny_cobyla
- initial_delta

Initial step size for non-gradient based optimizers

Specification

Alias: none
Argument(s): REAL

Description

If `initial_delta` is supplied by the user, it will be applied in an absolute sense in all coordinate directions. APPS documentation advocates choosing `initial_delta` to be the approximate distance from the initial point to the solution. If this is unknown, it is advisable to err on the side of choosing an `initial_delta` that is too large or to not specify it. In the latter case, APPS will take a full step to the boundary in each direction. Relative application of `initial_delta` is not available unless the user scales the problem accordingly.

threshold_delta

- Keywords Area
- method
- coliny_cobyla
- threshold_delta

Stopping criteria based on step length or pattern size

Specification

Alias: none
Argument(s): REAL

Description

`threshold_delta` is the step length or pattern size used to determine convergence.

solution_target

- Keywords Area
- method
- coliny_cobyla
- solution_target

Stopping criteria based on objective function value

Topics

This keyword is related to the topics:

- problem

Specification

Alias: solution_accuracy
Argument(s): REAL

Description

`solution_target` is a termination criterion. The algorithm will terminate when the function value falls below solution_target.

seed

- Keywords Area
- method
- coliny_cobyla
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

show_misc_options

- Keywords Area
- method
- coliny_cobyla
- show_misc_options

Show algorithm parameters not exposed in Dakota input

Specification

Alias: none
Argument(s): none

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

misc_options

- Keywords Area
- method
- coliny_cobyla
- misc_options

Set method options not available through Dakota spec

Specification

Alias: none
Argument(s): STRINGLIST

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

model_pointer

- Keywords Area
- method
- coliny_cobyla
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.42 coliny_direct

- Keywords Area
- method
- coliny_direct

DIviding RECTangles method

Topics

This keyword is related to the topics:

- package_scolib
- package_coliny
- global_optimization_methods

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		division	Determine how rectangles are subdivided
	Optional		global_balance_-parameter	Tolerance for whether a subregion is worth dividing
	Optional		local_balance_-parameter	Tolerance for whether a subregion is worth dividing
	Optional		max_boxsize_limit	Stopping Criterion based on longest edge of hyperrectangle
	Optional		min_boxsize_limit	Stopping Criterion based on shortest edge of hyperrectangle
	Optional		constraint_penalty	Multiplier for the penalty function
	Optional		solution_target	Stopping criteria based on objective function value
	Optional		seed	Specify the seed of the random number generator
	Optional		show_misc_options	Show algorithm parameters not exposed in Dakota input
	Optional		misc_options	Set method options not available through Dakota spec
	Optional		model_pointer	Specify which model block will be used by a method

Description

The DIviding RECTangles (DIRECT) optimization algorithm is a derivative free global optimization method that balances local search in promising regions of the design space with global search in unexplored regions. As shown in Figure 5.1, DIRECT adaptively subdivides the space of feasible design points so as to guarantee that iterates are generated in the neighborhood of a global minimum in finitely many iterations.

DIRECT" \image latex direct1.eps "Design space partitioning with DIRECT" width=10cm

In practice, DIRECT has proven an effective heuristic for engineering design applications, for which it is able

to quickly identify candidate solutions that can be further refined with fast local optimizers.

See the page package_scolib for important information regarding all SCOLIB methods

The DIRECT algorithm supports concurrency up to twice the number of variables being optimized.

DIRECT uses the `solution_target`, `constraint_penalty` and `show_misc_options` specifications that are described in package_scolib. Note, however, that DIRECT uses a fixed penalty value for constraint violations (i.e. it is not dynamically adapted as is done in `coliny_pattern_search`).

Search Parameters

The `global_balance_parameter` controls how much global search is performed by only allowing a subregion to be subdivided if the size of the subregion divided by the size of the largest subregion is at least `global_balance_parameter`. Intuitively, this forces large subregions to be subdivided before the smallest subregions are refined. The `local_balance_parameter` provides a tolerance for estimating whether the smallest subregion can provide a sufficient decrease to be worth subdividing; the default value is a small value that is suitable for most applications.

Stopping Critieria

DIRECT can be terminated with:

- max_function_evaluations
- max_iterations
- convergence_tolerance
- `solution_target`
- `max_boxsize_limit`
- `min_boxsize_limit` - most effective in practice

See Also

These keywords may also be of interest:

- coliny_beta
- coliny_pattern_search
- coliny_cobyla
- coliny_ea
- coliny_solis_wets
- ncsu_direct

division

- Keywords Area
- method
- coliny_direct
- division

Determine how rectangles are subdivided

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	major_dimension	(default) Longest edge of subregion is subdivided
			all_dimensions	All dimensions are simultaneously subdivided

Description

The `division` specification determines how DIRECT subdivides each subregion of the search space.

If `division` is set to `major_dimension`, then the dimension representing the longest edge of the subregion is subdivided (this is the default). If `division` is set to `all_dimensions`, then all dimensions are simultaneously subdivided.

major_dimension

- Keywords Area
- method
- coliny_direct
- division
- major_dimension

(default) Longest edge of subregion is subdivided

Specification

Alias: none

Argument(s): none

Description

Longest edge of subregion is subdivided

all_dimensions

- Keywords Area
- method
- coliny_direct
- division
- all_dimensions

All dimensions are simultaneously subdivided

Specification

Alias: none
Argument(s): none

Description

All dimensions are simultaneously subdivided

global_balance_parameter

- Keywords Area
- method
- coliny_direct
- global_balance_parameter

Tolerance for whether a subregion is worth dividing

Specification

Alias: none
Argument(s): REAL

Description

The `global_balance_parameter` controls how much global search is performed by only allowing a subregion to be subdivided if the size of the subregion divided by the size of the largest subregion is at least `global_balance_parameter`. Intuitively, this forces large subregions to be subdivided before the smallest subregions are refined.

local_balance_parameter

- Keywords Area
- method
- coliny_direct
- local_balance_parameter

Tolerance for whether a subregion is worth dividing

Specification

Alias: none
Argument(s): REAL

Description

See parent page. The `local_balance_parameter` provides a tolerance for estimating whether the smallest subregion can provide a sufficient decrease to be worth subdividing; the default value is a small value that is suitable for most applications.

max_boxsize_limit

- Keywords Area
- method
- coliny_direct
- max_boxsize_limit

Stopping Criterion based on longest edge of hyperrectangle

Specification

Alias: none
Argument(s): REAL

Description

Each subregion considered by DIRECT has a **size**, which corresponds to the longest diagonal of the subregion.

`max_boxsize_limit` specification terminates DIRECT if the size of the largest subregion falls below this threshold.

min_boxsize_limit

- Keywords Area
- method
- coliny_direct
- min_boxsize_limit

Stopping Criterion based on shortest edge of hyperrectangle

Specification

Alias: none
Argument(s): REAL

Description

`min_boxsize_limit` is a setting that terminates the optimization when the measure of a hyperrectangle S with f(c(S)) = fmin is less than min_boxsize_limit.

Each subregion considered by DIRECT has a **size**, which corresponds to the longest diagonal of the subregion.

`min_boxsize_limit` specification terminates DIRECT if the size of the smallest subregion falls below this threshold.

In practice, this specification is likely to be more effective at limiting DIRECT's search.

constraint_penalty

- Keywords Area
- method
- coliny_direct
- constraint_penalty

Multiplier for the penalty function

Specification

Alias: none
Argument(s): REAL

Description

- `constraint_penalty`: the initial multiplier for the penalty function, must be nonnegative

Most SCOLIB optimizers treat constraints with a simple penalty scheme that adds `constraint_penalty` times the sum of squares of the constraint violations to the objective function. Specific exceptions to this method for handling constraint violations are noted below. (The default value of `constraint_penalty` is 1000.0, except for methods that dynamically adapt their constraint penalty, for which the default value is 1.0.)

solution_target

- Keywords Area
- method
- coliny_direct
- solution_target

Stopping criteria based on objective function value

Topics

This keyword is related to the topics:

- problem

Specification

Alias: solution_accuracy
Argument(s): REAL

Description

`solution_target` is a termination criterion. The algorithm will terminate when the function value falls below solution_target.

seed

- Keywords Area
- method
- coliny_direct
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

show_misc_options

- Keywords Area
- method
- coliny_direct
- show_misc_options

Show algorithm parameters not exposed in Dakota input

Specification

Alias: none
Argument(s): none

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

misc_options

- Keywords Area
- method
- coliny_direct
- misc_options

Set method options not available through Dakota spec

Specification

Alias: none
Argument(s): STRINGLIST

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

model_pointer

- Keywords Area
- method
- coliny_direct
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.43 coliny_ea

- Keywords Area
- method
- coliny_ea

Evolutionary Algorithm

Topics

This keyword is related to the topics:

- package_scolib
- package_coliny
- global_optimization_methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		population_size	Set the initial population size in JEGA methods
	Optional		initialization_type	Specify how to initialize the population
	Optional		fitness_type	Select the fitness type for JEGA methods
	Optional		replacement_type	Select a replacement type for JEGA methods
	Optional		crossover_rate	Specify the probability of a crossover event

	Optional	crossover_type	Select a crossover type for JEGA methods
	Optional	mutation_rate	Set probability of a mutation
	Optional	mutation_type	Select a mutation type for JEGA methods
	Optional	constraint_penalty	Multiplier for the penalty function
	Optional	solution_target	Stopping criteria based on objective function value
	Optional	seed	Specify the seed of the random number generator
	Optional	show_misc_options	Show algorithm parameters not exposed in Dakota input
	Optional	misc_options	Set method options not available through Dakota spec
	Optional	model_pointer	Specify which model block will be used by a method

Description

Evolutionary Algorithm

See the page package_scolib for important information regarding all SCOLIB methods

`coliny_pattern_search` supports concurrency up to the size of the population

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results. The `population_size` control specifies how many individuals will comprise the EA's population.

The `initialization_type` defines the type of initialization for the population of the EA. There are three types: `simple_random`, `unique_random`, and `flat_file`. `simple_random` creates initial solutions with random variable values according to a uniform random number distribution. It gives no consideration to any previously generated designs. The number of designs is specified by the `population_size`. `unique_random` is the same as `simple_random`, except that when a new solution is generated, it is checked against the rest of the solutions. If it duplicates any of them, it is rejected. `flat_file` allows the initial population to be read from a flat file. If `flat_file` is specified, a file name must be given.

The `fitness_type` controls how strongly differences in "fitness" (i.e., the objective function) are weighted in the process of selecting "parents" for crossover:

- the `linear_rank` setting uses a linear scaling of probability of selection based on the rank order of each individual's objective function within the population

- the `merit_function` setting uses a proportional scaling of probability of selection based on the relative value of each individual's objective function within the population

The `replacement_type` controls how current populations and newly generated individuals are combined to create a new population. Each of the `replacement_type` selections accepts an integer value, which is referred to below as the `replacement_size`.

- The `random` setting creates a new population using (a) `replacement_size` randomly selected individuals from the current population, and (b) `population_size` - `replacement_size` individuals randomly selected from among the newly generated individuals (the number of which is optionally specified using `new_solutions_generated`) that are created for each generation (using the selection, crossover, and mutation procedures).

- The `chc` setting creates a new population using (a) the `replacement_size` best individuals from the *combination* of the current population and the newly generated individuals, and (b) `population_size` - `replacement_size` individuals randomly selected from among the remaining individuals in this combined pool. The `chc` setting is the preferred selection for many engineering problems.

- The `elitist` (default) setting creates a new population using (a) the `replacement_size` best individuals from the current population, (b) and `population_size` - `replacement_size` individuals randomly selected from the newly generated individuals. It is possible in this case to lose a good solution from the newly generated individuals if it is not randomly selected for replacement; however, the default `new_solutions_generated` value is set such that the entire set of newly generated individuals will be selected for replacement.

Note that `new_solutions_generated` is not recognized by Dakota as a valid keyword unless `replacement_type` has been specified.

Theory

The basic steps of an evolutionary algorithm are depicted in Figure 5.2.

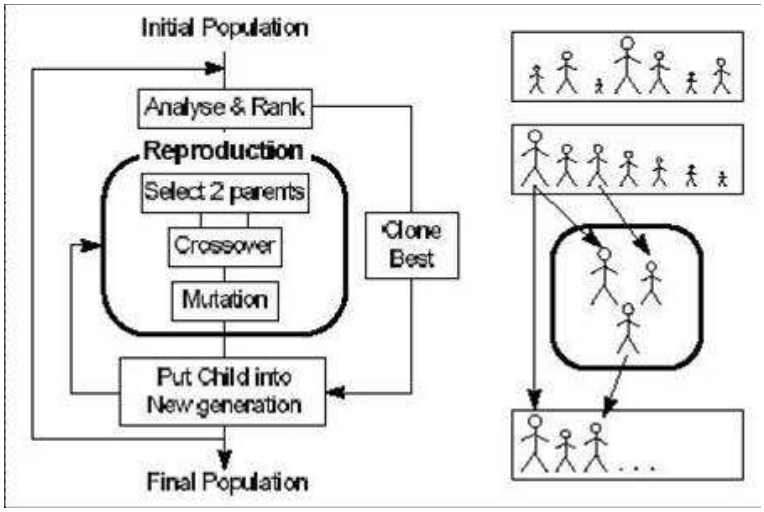

Figure 5.2: Depiction of evolutionary algorithm

They can be enumerated as follows:

1. Select an initial population randomly and perform function evaluations on these individuals
2. Perform selection for parents based on relative fitness
3. Apply crossover and mutation to generate `new_solutions_generated` new individuals from the selected parents
   - Apply crossover with a fixed probability from two selected parents
   - If crossover is applied, apply mutation to the newly generated individual with a fixed probability
   - If crossover is not applied, apply mutation with a fixed probability to a single selected parent
4. Perform function evaluations on the new individuals
5. Perform replacement to determine the new population
6. Return to step 2 and continue the algorithm until convergence criteria are satisfied or iteration limits are exceeded

See Also

These keywords may also be of interest:

- coliny_beta
- coliny_direct
- coliny_pattern_search
- coliny_cobyla
- coliny_solis_wets

population_size

- Keywords Area
- method
- coliny_ea
- population_size

Set the initial population size in JEGA methods

Specification

Alias: none
Argument(s): INTEGER

Description

The number of designs in the initial population is specified by the `population_size`. Note that the `population_size` only sets the size of the initial population. The population size may vary in the JEGA methods according to the type of operators chosen for a particular optimization run.

initialization_type

- Keywords Area
- method
- coliny_ea
- initialization_type

Specify how to initialize the population

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	simple_random	Initilization option for JEGA methods
			unique_random	(Default) Initilization option for JEGA methods
			flat_file	Initilization option for JEGA methods

Description

The `initialization_type` defines how the initial population is created for the GA. There are three types:

1. `simple_random`
2. `unique_random` (default)
3. `flat_file`

Setting the size for the `flat_file` initializer has the effect of requiring a minimum number of designs to create. If this minimum number has not been created once the files are all read, the rest are created using the `unique_random` initializer and then the `simple_random` initializer if necessary.

simple_random

- Keywords Area
- method
- coliny_ea
- initialization_type
- simple_random

Initilization option for JEGA methods

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

`simple_random` creates initial solutions with random variable values according to a uniform random number distribution. It gives no consideration to any previously generated designs.

unique_random

- Keywords Area
- method
- coliny_ea
- initialization_type
- unique_random

(Default) Initilization option for JEGA methods

Specification

Alias: none
Argument(s): none

Description

`unique_random` is the same as `simple_random`, except that when a new solution is generated, it is checked against the rest of the solutions. If it duplicates any of them, it is rejected.

flat_file

- Keywords Area
- method
- coliny_ea
- initialization_type
- flat_file

Initilization option for JEGA methods

Specification

Alias: none
Argument(s): STRING

Description

`flat_file` allows the initial population to be read from a flat file. If `flat_file` is specified, a file name must be given.

Variables can be delimited in the flat file in any way you see fit with a few exceptions. The delimiter must be the same on any given line of input with the exception of leading and trailing whitespace. So a line could look like: 1.1, 2.2 ,3.3 for example but could not look like: 1.1, 2.2 3.3. The delimiter can vary from line to line within the file which can be useful if data from multiple sources is pasted into the same input file. The delimiter can be any string that does not contain any of the characters .+-dDeE or any of the digits 0-9. The input will be read until the end of the file. The algorithm will discard any configurations for which it was unable to retrieve at least the number of design variables. The objective and constraint entries are not required but if ALL are present, they will be recorded and the design will be tagged as evaluated so that evaluators may choose not to re-evaluate them.

Setting the size for this initializer has the effect of requiring a minimum number of designs to create. If this minimum number has not been created once the files are all read, the rest are created using the `unique_random` initializer and then the `simple_random` initializer if necessary.

fitness_type

- Keywords Area
- method
- coliny_ea
- fitness_type

Select the fitness type for JEGA methods

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	linear_rank	Set selection scaling
			merit_function	Balance goals of reducing objective function and satisfying constraints

Description

The two JEGA methods use different fitness types, which are described on their respective pages.

linear_rank

- Keywords Area
- method
- coliny_ea
- fitness_type
- linear_rank

Set selection scaling

Specification

Alias: none
Argument(s): none

Description

The `linear_rank` setting uses a linear scaling of probability of selection based on the rank order of each individual's objective function within the population

merit_function

- Keywords Area
- method
- coliny_ea
- fitness_type
- merit_function

Balance goals of reducing objective function and satisfying constraints

Specification

Alias: none
Argument(s): none

Description

A `merit_function` is a function in constrained optimization that attempts to provide joint progress toward reducing the objective function and satisfying the constraints.

replacement_type

- Keywords Area
- method
- coliny_ea
- replacement_type

Select a replacement type for JEGA methods

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	random	Purely random samples
			chc	Create new population using replacement
			elitist	Use the best designs to form a new population
	Optional		new_solutions_-generated	Replace population with individuals chosen from population

Description

!!!!!!!!!! INCONSISTENCY in the manual. This says replacement types are the same between methods, however, this is not true in the input spec. below_limit vs. favor_feasible.

As of JEGA v2.0, all replacement types are common to both MOGA and SOGA. They include the `roulette_wheel`, `unique_roulette_wheel`, `elitist`, and `below_limit` selectors. In roulette_wheel replacement, each design is conceptually allotted a portion of a wheel proportional to its fitness relative to the fitnesses of the other Designs. Then, portions of the wheel are chosen at random and the design occupying those portions are duplicated into the next population. Those Designs allotted larger portions of the wheel are more likely to be selected (potentially many times). `unique_roulette_wheel` replacement is the same as `roulette_wheel` replacement, with the exception that a design may only be selected once. The `below_limit` selector attempts to keep all designs for which the negated fitness is below a certain limit. The values are negated to keep with the convention that higher fitness is better. The inputs to the `below_limit` selector are the limit as a real value, and a `shrinkage_percentage` as a real value. The `shrinkage_percentage` defines the minimum amount of selections that will take place if enough designs are available. It is interpreted as a percentage of the population size that must go on to the subsequent generation. To enforce this, `below_limit` makes all the selections it would make anyway and if that is not enough, it takes the remaining that it needs from the best of what is left (effectively raising its limit as far as it must to get the minimum number of selections). It continues until it has made enough selections. The `shrinkage_percentage` is designed to prevent extreme decreases in the population size at any given generation, and thus prevent a big loss of genetic diversity in a very short time. Without a shrinkage limit, a small group of "super" designs may appear and quickly cull the population down to a size on the order of the limiting value. In this case, all the diversity of the population is lost and it is expensive to re-diversify and spread the population. The `elitist` selector simply chooses the required number of designs

taking the most fit. For example, if 100 selections are requested, then the top 100 designs as ranked by fitness will be selected and the remaining will be discarded.

The initialization, crossover, and mutation controls were all described above. There are no SOGA specific aspects to these controls. The `replacement_type` for a SOGA may be `roulette_wheel`, `unique_roulette_wheel`, `elitist`, or `favor_feasible`.

The `favor_feasible` replacement type first considers feasibility as a selection criteria. If that does not produce a "winner" then it moves on to considering fitness value. Because of this, any fitness assessor used with the `favor_feasible` selector must only account objectives in the creation of fitness. Therefore, there is such a fitness assessor and it's use is enforced when the \ favor_feasible selector is chosen. In that case, and if the output level is set high enough, a message will be presented indicating that the `weighted_sum_only` fitness assessor will be used. As of JEGA v2.0 and beyond, the fitness assessment operator must be specified with SOGA although the `merit_function` is currently the only one (note that the `weighted_sum_only` assessor exists but cannot be selected). The roulette wheel selectors no longer assume a fitness function. The `merit_function` fitness assessor uses an exterior penalty function formulation to penalize infeasible designs. The specification allows the input of a `constraint_penalty` which is the multiplier to use on the constraint violations.

random

- Keywords Area
- method
- coliny_ea
- replacement_type
- random

Purely random samples

Specification

Alias: none
Argument(s): INTEGER

Description

Purely random samples.

chc

- Keywords Area
- method
- coliny_ea
- replacement_type
- chc

Create new population using replacement

Specification

Alias: none
Argument(s): INTEGER

Description

The `replacement_type` controls how current populations and newly generated individuals are combined to create a new population. Each of the `replacement_type` selections accepts an integer value, which is referred as the `replacement_size`:

The `chc` setting creates a new population using (a) the `replacement_size` best individuals from the *combination* of the current population and the newly generated individuals, and (b) `population_size` - `replacement_size` individuals randomly selected from among the remaining individuals in this combined pool. The `chc` setting is the preferred selection for many engineering problems.

elitist

- Keywords Area
- method
- coliny_ea
- replacement_type
- elitist

Use the best designs to form a new population

Specification

Alias: none
Argument(s): INTEGER

Description

The `elitist` (default) setting creates a new population using (a) the `replacement_size` best individuals from the current population, (b) and `population_size` - `replacement_size` individuals randomly selected from the newly generated individuals. It is possible in this case to lose a good solution from the newly generated individuals if it is not randomly selected for replacement; however, the default `new_solutions_generated` value is set such that the entire set of newly generated individuals will be selected for replacement.

new_solutions_generated

- Keywords Area
- method
- coliny_ea
- replacement_type
- new_solutions_generated

Replace population with individuals chosen from population

Specification

Alias: none
Argument(s): INTEGER

Description

- The `random` setting creates a new population using (a) `replacement_size` randomly selected individuals from the current population, and (b) `population_size` - `replacement_size` individuals randomly selected from among the newly generated individuals (the number of which is optionally specified using `new_solutions_generated`) that are created for each generation (using the selection, crossover, and mutation procedures).

crossover_rate

- Keywords Area
- method
- coliny_ea
- crossover_rate

Specify the probability of a crossover event

Specification

Alias: none
Argument(s): REAL

Description

The `crossover_type` controls what approach is employed for combining parent genetic information to create offspring, and the `crossover_rate` specifies the probability of a crossover operation being performed to generate a new offspring. The SCOLIB EA method supports three forms of crossover, `two_point`, `blend`, and `uniform`, which generate a new individual through combinations of two parent individuals. Two-point crossover divides each parent into three regions, where offspring are created from the combination of the middle region from one parent and the end regions from the other parent. Since the SCOLIB EA does not utilize bit representations of variable values, the crossover points only occur on coordinate boundaries, never within the bits of a particular coordinate. Uniform crossover creates offspring through random combination of coordinates from the two parents. Blend crossover generates a new individual randomly along the multidimensional vector connecting the two parents.

crossover_type

- Keywords Area
- method
- coliny_ea
- crossover_type

Select a crossover type for JEGA methods

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	Description of Group **Group 1**	Dakota Keyword two_point	Dakota Keyword Description Combine middle of one parent with end of another
			blend	Perform crossover by combining parents
			uniform	Set crossover type to be uniform

Description

There are many crossover types available. `multi_point_binary` crossover requires an integer number, N, of crossover points. This crossover type performs a bit switching crossover at N crossover points in the binary encoded genome of two designs. Thus, crossover may occur at any point along a solution chromosome (in the middle of a gene representing a design variable, for example). `multi_point_parameterized_binary` crossover is similar in that it performs a bit switching crossover routine at N crossover points. However, this crossover type performs crossover on each design variable individually. So the individual chromosomes are crossed at N locations. `multi_point_real` crossover performs a variable switching crossover routing at N crossover points in the real real valued genome of two designs. In this scheme, crossover only occurs between design variables (chromosomes). Note that the standard solution chromosome representation in the JEGA algorithm is real encoded and can handle integer or real design variables. For any crossover types that use a binary representation, real variables are converted to long integers by multiplying the real number by 10^6 and then truncating. Note that this assumes a precision of only six decimal places. Discrete variables are represented as integers (indices within a list of possible values) within the algorithm and thus require no special treatment by the binary operators.

The final crossover type is `shuffle_random`. This crossover type performs crossover by choosing design variables at random from a specified number of parents enough times that the requested number of children are produced. For example, consider the case of 3 parents producing 2 children. This operator would go through and for each design variable, select one of the parents as the donor for the child. So it creates a random shuffle of the parent design variable values. The relative numbers of children and parents are controllable to allow for as much mixing as desired. The more parents involved, the less likely that the children will wind up exact duplicates of the parents.

All crossover types take a `crossover_rate`. The crossover rate is used to calculate the number of crossover operations that take place. The number of crossovers is equal to the rate $*$ population_size.

two_point

- Keywords Area
- method
- coliny_ea
- crossover_type
- two_point

Combine middle of one parent with end of another

Specification

Alias: none
Argument(s): none

Description

The `crossover_type` controls what approach is employed for combining parent genetic information to create offspring, and the `crossover_rate` specifies the probability of a crossover operation being performed to generate a new offspring. The SCOLIB EA method supports three forms of crossover, `two_point`, `blend`, and `uniform`, which generate a new individual through combinations of two parent individuals. Two-point crossover divides each parent into three regions, where offspring are created from the combination of the middle region from one parent and the end regions from the other parent. Since the SCOLIB EA does not utilize bit representations of variable values, the crossover points only occur on coordinate boundaries, never within the bits of a particular coordinate. Uniform crossover creates offspring through random combination of coordinates from the two parents. Blend crossover generates a new individual randomly along the multidimensional vector connecting the two parents.

blend

- Keywords Area
- method
- coliny_ea
- crossover_type
- blend

Perform crossover by combining parents

Specification

Alias: none
Argument(s): none

Description

The `crossover_type` controls what approach is employed for combining parent genetic information to create offspring, and the `crossover_rate` specifies the probability of a crossover operation being performed to generate a new offspring. The SCOLIB EA method supports three forms of crossover, `two_point`, `blend`, and `uniform`, which generate a new individual through combinations of two parent individuals. Two-point crossover divides each parent into three regions, where offspring are created from the combination of the middle region from one parent and the end regions from the other parent. Since the SCOLIB EA does not utilize bit representations of variable values, the crossover points only occur on coordinate boundaries, never within the bits of a particular coordinate. Uniform crossover creates offspring through random combination of coordinates from the two parents. Blend crossover generates a new individual randomly along the multidimensional vector connecting the two parents.

uniform

- Keywords Area
- method
- coliny_ea
- crossover_type
- uniform

Set crossover type to be uniform

Specification

Alias: none
Argument(s): none

Description

The `crossover_type` controls what approach is employed for combining parent genetic information to create offspring, and the `crossover_rate` specifies the probability of a crossover operation being performed to generate a new offspring. The SCOLIB EA method supports three forms of crossover, `two_point`, `blend`, and `uniform`, which generate a new individual through combinations of two parent individuals. Two-point crossover divides each parent into three regions, where offspring are created from the combination of the middle region from one parent and the end regions from the other parent. Since the SCOLIB EA does not utilize bit representations of variable values, the crossover points only occur on coordinate boundaries, never within the bits of a particular coordinate. Uniform crossover creates offspring through random combination of coordinates from the two parents. Blend crossover generates a new individual randomly along the multidimensional vector connecting the two parents.

mutation_rate

- Keywords Area
- method
- coliny_ea
- mutation_rate

Set probability of a mutation

Specification

Alias: none
Argument(s): REAL

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

mutation_type

- Keywords Area
- method
- coliny_ea
- mutation_type

Select a mutation type for JEGA methods

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** replace_uniform	**Dakota Keyword Description** Use uniformly distributed value over range of parameter
			offset_normal	Set mutation offset to use a normal distribution

			offset_cauchy	Use a Cauchy distribution for the mutation offset
			offset_uniform	Set mutation offset to use a uniform distribution
	Optional		non_adaptive	

Description

There are five mutation types allowed. `replace_uniform` introduces random variation by first randomly choosing a design variable of a randomly selected design and reassigning it to a random valid value for that variable. No consideration of the current value is given when determining the new value. All mutation types have a `mutation_rate`. The number of mutations for the replace_uniform mutator is the product of the mutation_rate and the population_size.

The `bit_random` mutator introduces random variation by first converting a randomly chosen variable of a randomly chosen design into a binary string. It then flips a randomly chosen bit in the string from a 1 to a 0 or visa versa. In this mutation scheme, the resulting value has more probability of being similar to the original value. The number of mutations performed is the product of the mutation_rate, the number of design variables, and the population size.

The offset mutators all act by adding an "offset" random amount to a variable value. The random amount has a mean of zero in all cases. The `offset_normal` mutator introduces random variation by adding a Gaussian random amount to a variable value. The random amount has a standard deviation dependent on the `mutation_scale`. The `mutation_scale` is a fraction in the range [0, 1] and is meant to help control the amount of variation that takes place when a variable is mutated. `mutation_scale` is multiplied by the range of the variable being mutated to serve as standard deviation. `offset_cauchy` is similar to `offset_normal`, except that a Cauchy random variable is added to the variable being mutated. The `mutation_scale` also defines the standard deviation for this mutator. Finally, `offset_uniform` adds a uniform random amount to the variable value. For the `offset_uniform` mutator, the `mutation_scale` is interpreted as a fraction of the total range of the variable. The range of possible deviation amounts is +/- 1/2 $*$ (`mutation_scale` $*$ variable range). The number of mutations for all offset mutators is defined as the product of `mutation_rate` and `population_size`.

replace_uniform

- Keywords Area
- method
- coliny_ea
- mutation_type
- replace_uniform

Use uniformly distributed value over range of parameter

Specification

Alias: none

Argument(s): none

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

offset_normal

- Keywords Area
- method
- coliny_ea
- mutation_type
- offset_normal

Set mutation offset to use a normal distribution

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		mutation_scale	Scales mutation across range of parameter
	Optional		mutation_range	Set uniform offset control for discrete parameters

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given

coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

mutation_scale

- Keywords Area
- method
- coliny_ea
- mutation_type
- offset_normal
- mutation_scale

Scales mutation across range of parameter

Specification

Alias: none
Argument(s): REAL

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These

offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

Note that mutation_scale and mutation_range are not recognized by Dakota as valid keywords unless mutation_type has been specified and the type is an "offset" mutations.

mutation_range

- Keywords Area
- method
- coliny_ea
- mutation_type
- offset_normal
- mutation_range

Set uniform offset control for discrete parameters

Specification

Alias: none

Argument(s): INTEGER

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

offset_cauchy

- Keywords Area
- method

- coliny_ea
- mutation_type
- offset_cauchy

Use a Cauchy distribution for the mutation offset

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		mutation_scale	Scales mutation across range of parameter
	Optional		mutation_range	Set uniform offset control for discrete parameters

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

mutation_scale

- Keywords Area
- method
- coliny_ea
- mutation_type
- offset_cauchy

- mutation_scale

Scales mutation across range of parameter

Specification

Alias: none
Argument(s): REAL

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

Note that `mutation_scale` and `mutation_range` are not recognized by Dakota as valid keywords unless `mutation_type` has been specified and the type is an "offset" mutations.

mutation_range

- Keywords Area
- method
- coliny_ea
- mutation_type
- offset_cauchy
- mutation_range

Set uniform offset control for discrete parameters

Specification

Alias: none
Argument(s): INTEGER

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

offset_uniform

- Keywords Area
- method
- coliny_ea
- mutation_type
- offset_uniform

Set mutation offset to use a uniform distribution

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		mutation_scale	Scales mutation across range of parameter
	Optional		mutation_range	Set uniform offset control for discrete parameters

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given

coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by `mutation_scale`. The `replace_uniform` mutation type is not limited by `mutation_scale`; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

mutation_scale

- Keywords Area
- method
- coliny_ea
- mutation_type
- offset_uniform
- mutation_scale

Scales mutation across range of parameter

Specification

Alias: none
Argument(s): REAL

Description

The `mutation_type` controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the `offset_uniform` method. The `mutation_rate` controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The `mutation_scale` specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so `mutation_scale` is a relative value between 0 and 1. The `mutation_range` is used to control `offset_uniform` mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to `mutation_range`. The `offset_normal`, `offset_cauchy`, and `offset_uniform` mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These

offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

Note that mutation_scale and mutation_range are not recognized by Dakota as valid keywords unless mutation_type has been specified and the type is an "offset" mutations.

mutation_range

- Keywords Area
- method
- coliny_ea
- mutation_type
- offset_uniform
- mutation_range

Set uniform offset control for discrete parameters

Specification

Alias: none

Argument(s): INTEGER

Description

The mutation_type controls what approach is employed in randomly modifying continuous design variables within the EA population. Each of the mutation methods generates coordinate-wise changes to individuals, usually by adding a random variable to a given coordinate value (an "offset" mutation), but also by replacing a given coordinate value with a random variable (a "replace" mutation). Discrete design variables are always mutated using the offset_uniform method. The mutation_rate controls the probability of mutation being performed on an individual, both for new individuals generated by crossover (if crossover occurs) and for individuals from the existing population. It is the fraction of trial points that are mutated in a given iteration and therefore must be specified to be between 0 and 1. When mutation is performed, all dimensions of each individual are mutated. The mutation_scale specifies a scale factor which scales continuous mutation offsets; this is a fraction of the total range of each dimension, so mutation_scale is a relative value between 0 and 1. The mutation_range is used to control offset_uniform mutation used for discrete parameters. The replacement discrete value is the original value plus or minus an integer value up to mutation_range. The offset_normal, offset_cauchy, and offset_uniform mutation types are "offset" mutations in that they add a 0-mean random variable with a normal, cauchy, or uniform distribution, respectively, to the existing coordinate value. These offsets are limited in magnitude by mutation_scale. The replace_uniform mutation type is not limited by mutation_scale; rather it generates a replacement value for a coordinate using a uniformly distributed value over the total range for that coordinate.

non_adaptive

- Keywords Area
- method

- coliny_ea
- mutation_type
- non_adaptive

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): none

constraint_penalty

- Keywords Area
- method
- coliny_ea
- constraint_penalty

Multiplier for the penalty function

Specification

Alias: none
Argument(s): REAL

Description

- `constraint_penalty`: the initial multiplier for the penalty function, must be nonnegative

Most SCOLIB optimizers treat constraints with a simple penalty scheme that adds `constraint_penalty` times the sum of squares of the constraint violations to the objective function. Specific exceptions to this method for handling constraint violations are noted below. (The default value of `constraint_penalty` is 1000.0, except for methods that dynamically adapt their constraint penalty, for which the default value is 1.0.)

solution_target

- Keywords Area
- method
- coliny_ea
- solution_target

Stopping criteria based on objective function value

Topics

This keyword is related to the topics:

- problem

Specification

Alias: solution_accuracy
Argument(s): REAL

Description

`solution_target` is a termination criterion. The algorithm will terminate when the function value falls below solution_target.

seed

- Keywords Area
- method
- coliny_ea
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

show_misc_options

- Keywords Area
- method
- coliny_ea
- show_misc_options

Show algorithm parameters not exposed in Dakota input

Specification

Alias: none
Argument(s): none

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

misc_options

- Keywords Area
- method
- coliny_ea
- misc_options

Set method options not available through Dakota spec

Specification

Alias: none
Argument(s): STRINGLIST

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

model_pointer

- Keywords Area
- method
- coliny_ea
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.44 coliny_beta

- Keywords Area
- method
- coliny_beta

What does this do??????????????

Topics

This keyword is related to the topics:

- package_scolib
- package_coliny
- problem
- not_yet_reviewed

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		beta_solver_name	Use an in-development SCOLIB solver

	Optional	solution_target	Stopping criteria based on objective function value
	Optional	seed	Specify the seed of the random number generator
	Optional	show_misc_options	Show algorithm parameters not exposed in Dakota input
	Optional	misc_options	Set method options not available through Dakota spec
	Optional	model_pointer	Specify which model block will be used by a method

Description

See the page package_scolib for important information
!!!!!! NOT ACTUALLY MENTIONED AT ALL IN THE REFMAN??!?!?!

See Also

These keywords may also be of interest:

- coliny_direct
- coliny_pattern_search
- coliny_cobyla
- coliny_ea
- coliny_solis_wets

beta_solver_name

- Keywords Area
- method
- coliny_beta
- beta_solver_name

Use an in-development SCOLIB solver

Specification

Alias: none
Argument(s): STRING

Description

This is a means of accessing new methods in SCOLIB before they are exposed through the Dakota interface. Seek help from a Dakota or SCOLIB developer or a Dakota developer.

solution_target

- Keywords Area
- method
- coliny_beta
- solution_target

Stopping criteria based on objective function value

Topics

This keyword is related to the topics:

- problem

Specification

Alias: solution_accuracy
Argument(s): REAL

Description

`solution_target` is a termination criterion. The algorithm will terminate when the function value falls below solution_target.

seed

- Keywords Area
- method
- coliny_beta
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

show_misc_options

- Keywords Area
- method
- coliny_beta
- show_misc_options

Show algorithm parameters not exposed in Dakota input

Specification

Alias: none
Argument(s): none

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

misc_options

- Keywords Area
- method
- coliny_beta
- misc_options

Set method options not available through Dakota spec

Specification

Alias: none
Argument(s): STRINGLIST

Description

All SCOLIB methods support the `show_misc_options` optional specification which results in a dump of all the allowable method inputs. Note that the information provided by this command refers to optimizer parameters that are internal to SCOLIB, and which may differ from corresponding parameters used by the Dakota interface. The `misc_options` optional specification provides a means for inputing additional settings supported by the SCOLIB methods but which are not currently mapped through the Dakota input specification. Care must be taken in using this specification; they should only be employed by users familiar with the full range of parameter specifications available directly from SCOLIB and understand any differences that exist between those specifications and the ones available through Dakota.

model_pointer

- Keywords Area
- method
- coliny_beta
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.45 nl2sol

- Keywords Area
- method
- nl2sol

Trust-region method for nonlinear least squares

Topics

This keyword is related to the topics:

- nonlinear_least_squares

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		function_precision	Specify the maximum precision of the analysis code responses
	Optional		absolute_conv_tol	Absolute convergence tolerance
	Optional		x_conv_tol	X-convergence tolerance
	Optional		singular_conv_tol	Singular convergence tolerance
	Optional		singular_radius	
	Optional		false_conv_tol	False convergence tolerance
	Optional		initial_trust_radius	Initial trust region radius
	Optional		covariance	Determine how the final covariance matrix is computed
	Optional		regression_-diagnostics	Turn on regression diagnostics
	Optional		model_pointer	Specify which model block will be used by a method

Description

NL2SOL is available as `nl2sol` and addresses unconstrained and bound-constrained least squares problems. It uses a trust-region method (and thus can be viewed as a generalization of the Levenberg-Marquardt algorithm) and adaptively chooses between two Hessian approximations, the Gauss-Newton approximation alone and the Gauss-Newton approximation plus a quasi-Newton approximation to the rest of the Hessian. Even on small-residual problems, the latter Hessian approximation can be useful when the starting guess is far from the solution. On problems that are not over-parameterized (i.e., that do not involve more optimization variables than the data support), NL2SOL usually exhibits fast convergence.

Several internal NL2SOL convergence tolerances are adjusted in response to `function_precision`, which gives the relative precision to which responses are computed.

These tolerances may also be specified explicitly using:

- `convergence_tolerance` (NL2SOL's `rfctol`)
- `x_conv_tol` (NL2SOL's `xctol`)
- `absolute_conv_tol` (NL2SOL's `afctol`)
- `singular_conv_tol` (NL2SOL's `sctol`)
- `false_conv_tol` (NL2SOL's `xftol`)

- initial_trust_radius (NL2SOL's lmax0)

The internal NL2SOL defaults can be obtained for many of these controls by specifying the value -1. The internal defaults are often functions of machine epsilon (as limited by function_precision).

Examples

An example of nl2sol is given below, and is discussed in the User's Manual.

Note that in this usage of calibration_terms, the driver script rosenbrock, is returning "residuals", which the nl2sol method is attempting to minimze. Another use case is to provide a data file, which Dakota will attempt to match the model responses to. See calibration_data_file.

```
# Dakota Input File: rosen_opt_nls.in
environment
  tabular_graphics_data
    tabular_graphics_file = 'rosen_opt_nls.dat'

method
  max_iterations = 100
  convergence_tolerance = 1e-4
  nl2sol

model
  single

variables
  continuous_design = 2
    initial_point    -1.2      1.0
    lower_bounds     -2.0     -2.0
    upper_bounds      2.0      2.0
    descriptors       'x1'     "x2"

interface
  analysis_driver = 'rosenbrock'
    direct

responses
  calibration_terms = 2
  analytic_gradients
  no_hessians
```

Theory

NL2SOL has a variety of internal controls as described in AT&T Bell Labs CS TR 153 (http://cm.bell-labs.-com/cm/cs/cstr/153.ps.gz). A number of existing Dakota controls (method independent controls and responses controls) are mapped into these NL2SOL internal controls. In particular, Dakota's convergence-_tolerance, max_iterations, max_function_evaluations, and fd_gradient_step_size are mapped directly into NL2SOL's rfctol, mxiter, mxfcal, and dltfdj controls, respectively. In addition, Dakota's fd_hessian_step_size is mapped into both delta0 and dltfdc, and Dakota's output verbosity is mapped into NL2SOL's auxprt and outlev (for normal/verbose/debug output, N-L2SOL prints initial guess, final solution, solution statistics, nondefault values, and changes to the active bound constraint set on every iteration; for quiet output, NL2SOL prints only the initial guess and final solution; and for silent output, NL2SOL output is suppressed).

See Also

These keywords may also be of interest:

- nlssol_sqp
- optpp_g_newton

function_precision

- Keywords Area
- method
- nl2sol
- function_precision

Specify the maximum precision of the analysis code responses

Specification

Alias: none
Argument(s): REAL

Description

The `function_precision` control provides the algorithm with an estimate of the accuracy to which the problem functions can be computed. This is used to prevent the algorithm from trying to distinguish between function values that differ by less than the inherent error in the calculation.

absolute_conv_tol

- Keywords Area
- method
- nl2sol
- absolute_conv_tol

Absolute convergence tolerance

Specification

Alias: none
Argument(s): REAL

Description

`absolute_conv_tol` (NL2SOL's `afctol`) is the absolute function convergence tolerance (stop when half the sum of squares is less than `absolute_conv_tol`, which is mainly of interest on zero-residual test problems).

The internal default is a function of machine epsilon (as limited by `function_precision`). The default is selected with a value of -1.

x_conv_tol

- Keywords Area
- method
- nl2sol
- x_conv_tol

X-convergence tolerance

Specification

Alias: none
Argument(s): REAL

Description

`x_conv_tol` maps to the internal NL2SOL control `xctol`. It is the X-convergence tolerance (scaled relative accuracy of the solution variables).

The internal default is a function of machine epsilon (as limited by `function_precision`). The default is selected with a value of -1.

singular_conv_tol

- Keywords Area
- method
- nl2sol
- singular_conv_tol

Singular convergence tolerance

Specification

Alias: none
Argument(s): REAL

Description

`singular_conv_tol` (NL2SOL's `sctol`) is the singular convergence tolerance, which works in conjunction with singular_radius to test for underdetermined least-squares problems (stop when the relative reduction yet possible in the sum of squares appears less then `singular_conv_tol` for steps of scaled length at most singular_radius).

The internal default is a function of machine epsilon (as limited by `function_precision`). The default is selected with a value of -1.

singular_radius

- Keywords Area
- method
- nl2sol
- singular_radius

TODO: Need a Blurb

Specification

Alias: none
Argument(s): REAL

Description

`singular_radius` works in conjunction with singular_conv_tol to test for underdetermined least-squares problems (stop when the relative reduction yet possible in the sum of squares appears less then `singular_conv_tol` for steps of scaled length at most singular_radius).

The internal default results in the internal use of steps of length 1. The default is selected with a value of -1.

false_conv_tol

- Keywords Area
- method
- nl2sol
- false_conv_tol

False convergence tolerance

Specification

Alias: none
Argument(s): REAL

Description

`false_conv_tol` (NL2SOL's `xftol`) is the false-convergence tolerance (stop with a suspicion of discontinuity when a more favorable stopping test is not satisfied and a step of scaled length at most `false_conv_tol` is not accepted)

The internal default is a function of machine epsilon (as limited by `function_precision`). The default is selected with a value of -1.

initial_trust_radius

- Keywords Area
- method
- nl2sol
- initial_trust_radius

Initial trust region radius

Specification

Alias: none
Argument(s): REAL

Description

`initial_trust_radius` specification (NL2SOL's `lmax0`) specifies the initial trust region radius for the algorithm.

The internal default results in the internal use of steps of length 1. The default is selected with a value of -1.

covariance

- Keywords Area
- method
- nl2sol
- covariance

Determine how the final covariance matrix is computed

Specification

Alias: none
Argument(s): INTEGER

Description

`covariance` (NL2SOL's `covreq`) specifies whether and how NL2SOL computes a final covariance matrix.

The desired covariance approximation:

- 0 = default = none
- 1 or -1 ==> $\sigma^2 H^{-1} J^T J H^{-1}$
- 2 or -2 ==> $\sigma^2 H^{-1}$
- 3 or -3 ==> $\sigma^2 (J^T J)^{-1}$
- Negative values ==> estimate the final Hessian H by finite differences of function values only (using `fd_hessian_step_size`)
- Positive values ==> differences of gradients (using `fd_hessian_step_size`)

regression_diagnostics

- Keywords Area
- method
- nl2sol
- regression_diagnostics

Turn on regression diagnostics

Specification

Alias: none
Argument(s): none

Description

When `regression_diagnostics` (NL2SOL's `rdreq`) is specified and a positive-definite final Hessian approximation H is computed, NL2SOL computes and prints a regression diagnostic vector RD such that if omitting the i-th observation would cause alpha times the change in the solution that omitting the j-th observation would cause, then RD[i] = |alpha| RD[j]. The finite-difference step-size tolerance affecting H is fd_step_size (NL2SOL's `delta0` and `dltfdc`).

model_pointer

- Keywords Area
- method
- nl2sol
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.46 nonlinear_cg

- Keywords Area
- method
- nonlinear_cg

This is not mentioned in the reference manual

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		misc_options	
	Optional		model_pointer	Specify which model block will be used by a method

misc_options

- Keywords Area
- method
- nonlinear_cg
- misc_options

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none

Argument(s): STRINGLIST

model_pointer

- Keywords Area
- method
- nonlinear_cg
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.47 ncsu_direct

- Keywords Area
- method
- ncsu_direct

DIviding RECTangles method

Topics

This keyword is related to the topics:

- global_optimization_methods

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		solution_target	Specifies a globally optimal value toward which the optimizer should track
	Optional		min_boxsize_limit	Stopping Criterion based on shortest edge of hyperrectangle
	Optional		volume_boxsize_-limit	Stopping criterion based on volume of search space
	Optional		model_pointer	Specify which model block will be used by a method

Description

North Carolina State University (NCSU) has an implementation of the DIRECT algorithm (DIviding RECTangles algorithm that is outlined in the SCOLIB method section above). This version is documented in [[27] "Gablonsky, 2001".] We have found that the NCSU DIRECT implementation works better and is more robust for some problems than `coliny_direct`. Currently, we maintain both versions of DIRECT in Dakota; in the future, we may deprecate one.

The NCSU DIRECT method is selected with `ncsu_direct`. We have tried to maintain consistency between the keywords in SCOLIB and NCSU implementation of DIRECT, but the algorithms have different parameters, so the keywords sometimes have slightly different meaning.

Stopping Criteria

The algorithm stops based on:

1. max_iterations - number of iterations
2. max_function_evaluations - number of function evaluations
3. `solution_target` and convergence_tolerance
4. `min_boxsize_limit`
5. `volume_boxsize_limit`

This method will always strictly respect the number of iterations, but may slightly exceed the number of function evaluations, as it will always explore all sub-rectangles at the current level.

See Also

These keywords may also be of interest:

- coliny_direct

solution_target

- Keywords Area
- method
- ncsu_direct
- solution_target

Specifies a globally optimal value toward which the optimizer should track

Specification

Alias: solution_accuracy
Argument(s): REAL

Description

The solution target specifies a goal toward which the optimizer should track.

This is used for test problems, when the true global minimum is known (call it `solution_target` := fglobal). Then, the optimization terminates when 100(f_min-fglobal)/max(1,abs(fglobal) < convergence_tolerance. The default for fglobal is -1.0e100 and the default for convergence tolerance is described at convergence_tolerance.

min_boxsize_limit

- Keywords Area
- method
- ncsu_direct
- min_boxsize_limit

Stopping Criterion based on shortest edge of hyperrectangle

Specification

Alias: none
Argument(s): REAL

Description

`min_boxsize_limit` is a setting that terminates the optimization when the measure of a hyperrectangle S with f(c(S)) = fmin is less than min_boxsize_limit.

Each subregion considered by DIRECT has a **size**, which corresponds to the longest diagonal of the subregion.

`min_boxsize_limit` specification terminates DIRECT if the size of the smallest subregion falls below this threshold.

In practice, this specification is likely to be more effective at limiting DIRECT's search.

volume_boxsize_limit

- Keywords Area
- method
- ncsu_direct
- volume_boxsize_limit

Stopping criterion based on volume of search space

Specification

Alias: none
Argument(s): REAL

Description

`volume_boxsize_limit` is a setting that terminates the optimization when the volume of a hyperrectangle S with f(c(S)) = fmin is less than volume_boxsize_limit percent of the original hyperrectangle. Basically, volume_boxsize_limit stops the optimization when the volume of the particular rectangle which has fmin is less than a certain percentage of the whole volume.

model_pointer

- Keywords Area
- method
- ncsu_direct
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.48 genie_opt_darts

- Keywords Area
- method
- genie_opt_darts

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		seed	Specify the seed of the random number generator
	Optional		model_pointer	Specify which model block will be used by a method

seed

- Keywords Area
- method
- genie_opt_darts
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- genie_opt_darts
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.49 genie_direct

- Keywords Area
- method
- genie_direct

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		seed	Specify the seed of the random number generator
	Optional		model_pointer	Specify which model block will be used by a method

seed

- Keywords Area
- method
- genie_direct
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- genie_direct
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.50 efficient_global

- Keywords Area
- method
- efficient_global

Global Surrogate Based Optimization, a.k.a. EGO

Topics

This keyword is related to the topics:

- global_optimization_methods
- surrogate_based_optimization_methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		gaussian_process	Use a Gaussian Process surrogate
	Optional		use_derivatives	Use derivative data to construct surrogate models
	Optional		import_points_file	Specify file containing points (variable and response values)
	Optional		export_points_file	Specify file to write out the responses from the surrogate
	Optional		seed	Specify the seed of the random number generator

	Optional	model_pointer	Specify which model block will be used by a method

Description

The Efficient Global Optimization (EGO) method was first developed by Jones, Schonlau, and Welch [[55] "Jones et al., 1998"]. In EGO, a stochastic response surface approximation for the objective function is developed based on some sample points from the "true" simulation.

Note that several major differences exist between our implementation and that of [[55] "Jones et al., 1998"]. First, rather than using a branch and bound method to find the point which maximizes the EIF, we use the DIRECT global optimization method.

Second, we support both global optimization and global nonlinear least squares as well as general nonlinear constraints through abstraction and subproblem recasting.

The efficient global method is in prototype form. Currently, we do not expose any specification controls for the underlying Gaussian process model used or for the optimization of the expected improvement function (which is currently performed by the NCSU DIRECT algorithm using its internal defaults).

By default, EGO uses the Surfpack GP (Kriging) model, but the Dakota implementation may be selected instead. If `use_derivatives` is specified the GP model will be built using available derivative data (Surfpack GP only).

Theory

The particular response surface used is a Gaussian process (GP). The GP allows one to calculate the prediction at a new input location as well as the uncertainty associated with that prediction. The key idea in EGO is to maximize the Expected Improvement Function (EIF). The EIF is used to select the location at which a new training point should be added to the Gaussian process model by maximizing the amount of improvement in the objective function that can be expected by adding that point. A point could be expected to produce an improvement in the objective function if its predicted value is better than the current best solution, or if the uncertainty in its prediction is such that the probability of it producing a better solution is high. Because the uncertainty is higher in regions of the design space with few observations, this provides a balance between exploiting areas of the design space that predict good solutions, and exploring areas where more information is needed. EGO trades off this "exploitation vs. exploration." The general procedure for these EGO-type methods is:

- Build an initial Gaussian process model of the objective function
- Find the point that maximizes the EIF. If the EIF value at this point is sufficiently small, stop.
- Evaluate the objective function at the point where the EIF is maximized. Update the Gaussian process model using this new point.
- Return to the previous step.

See Also

These keywords may also be of interest:

- surrogate_based_local
- surrogate_based_global

gaussian_process

- Keywords Area
- method
- efficient_global
- gaussian_process

Use a Gaussian Process surrogate

Specification

Alias: kriging

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	surfpack	Use the Surfpack version of Gaussian Process surrogates
			dakota	Select the built in Gaussian Process surrogate

Description

Dakota uses the Gaussian process (GP) surrogate from Surfpack, which is specified using the surfpack keyword.

A second version of GP surrogates was available in prior versions of Dakota.

For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.

surfpack

- Keywords Area
- method
- efficient_global
- gaussian_process
- surfpack

Use the Surfpack version of Gaussian Process surrogates

Specification

Alias: none

Argument(s): none

Description

This keyword specifies the use of the Gaussian process that is incorporated in our surface fitting library called Surfpack.

Several user options are available:

1. Optimization methods:

   Maximum Likelihood Estimation (MLE) is used to find the optimal values of the hyper-parameters governing the trend and correlation functions. By default the global optimization method DIRECT is used for MLE, but other options for the optimization method are available. See optimization_method.

   The total number of evaluations of the likelihood function can be controlled using the `max_trials` keyword followed by a positive integer. Note that the likelihood function does not require running the "truth" model, and is relatively inexpensive to compute.

2. Trend Function:

   The GP models incorporate a parametric trend function whose purpose is to capture large-scale variations. See trend.

3. Correlation Lengths:

   Correlation lengths are usually optimized by Surfpack, however, the user can specify the lengths manually. See correlation_lengths.

4. Ill-conditioning

   One of the major problems in determining the governing values for a Gaussian process or Kriging model is the fact that the correlation matrix can easily become ill-conditioned when there are too many input points close together. Since the predictions from the Gaussian process model involve inverting the correlation matrix, ill-conditioning can lead to poor predictive capability and should be avoided.

   Note that a sufficiently bad sample design could require correlation lengths to be so short that any interpolatory GP model would become inept at extrapolation and interpolation.

   The `surfpack` model handles ill-conditioning internally by default, but behavior can be modified using

5. Gradient Enhanced Kriging (GEK).

   The `use_derivatives` keyword will cause the Surfpack GP to be constructed from a combination of function value and gradient information (if available).

   See notes in the Theory section.

Theory

Gradient Enhanced Kriging

Incorporating gradient information will only be beneficial if accurate and inexpensive derivative information is available, and the derivatives are not infinite or nearly so. Here "inexpensive" means that the cost of evaluating a function value plus gradient is comparable to the cost of evaluating only the function value, for example gradients computed by analytical, automatic differentiation, or continuous adjoint techniques. It is not cost effective to use derivatives computed by finite differences. In tests, GEK models built from finite difference derivatives were also significantly less accurate than those built from analytical derivatives. Note that GEK's correlation matrix tends to have a significantly worse condition number than Kriging for the same sample design.

This issue was addressed by using a pivoted Cholesky factorization of Kriging's correlation matrix (which is a small sub-matrix within GEK's correlation matrix) to rank points by how much unique information they contain. This reordering is then applied to whole points (the function value at a point immediately followed by gradient

information at the same point) in GEK's correlation matrix. A standard non-pivoted Cholesky is then applied to the reordered GEK correlation matrix and a bisection search is used to find the last equation that meets the constraint on the (estimate of) condition number. The cost of performing pivoted Cholesky on Kriging's correlation matrix is usually negligible compared to the cost of the non-pivoted Cholesky factorization of GEK's correlation matrix. In tests, it also resulted in more accurate GEK models than when pivoted Cholesky or whole-point-block pivoted Cholesky was performed on GEK's correlation matrix.

?? This correlation function results in a response surface that is C^{∞}-continuous.

dakota

- Keywords Area
- method
- efficient_global
- gaussian_process
- dakota

Select the built in Gaussian Process surrogate

Specification

Alias: none

Argument(s): none

Description

A second version of GP surrogates was available in prior versions of Dakota. **For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.**

Historically these models were drastically different, but in Dakota 5.1, they became quite similar. They now differ in that the Surfpack GP has a richer set of features/options and tends to be more accurate than the Dakota version. Due to how the Surfpack GP handles ill-conditioned correlation matrices (which significantly contributes to its greater accuracy), the `Surfpack` GP can be a factor of two or three slower than Dakota's. As of Dakota 5.2, the Surfpack implementation is the default in all contexts except Bayesian calibration.

More details on the `gaussian_process` dakota model can be found in[59].

Dakota's GP deals with ill-conditioning in two ways. First, when it encounters a non-invertible correlation matrix it iteratively increases the size of a "nugget," but in such cases the resulting approximation smooths rather than interpolates the data. Second, it has a `point_selection` option (default off) that uses a greedy algorithm to select a well-spaced subset of points prior to the construction of the GP. In this case, the GP will only interpolate the selected subset. Typically, one should not need point selection in trust-region methods because a small number of points are used to develop a surrogate within each trust region. Point selection is most beneficial when constructing with a large number of points, typically more than order one hundred, though this depends on the number of variables and spacing of the sample points.

This differs from the `point_selection` option of the Dakota GP which initially chooses a well-spaced subset of points and finds the correlation parameters that are most likely for that one subset.

use_derivatives

- Keywords Area
- method
- efficient_global
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none
Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

import_points_file

- Keywords Area
- method
- efficient_global
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- efficient_global
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- efficient_global
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- efficient_global
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- efficient_global
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- efficient_global
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

seed

- Keywords Area
- method
- efficient_global
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- efficient_global
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.51 polynomial_chaos

- Keywords Area
- method
- polynomial_chaos

Uncertainty quantification using polynomial chaos expansions

Specification

Alias: nond_polynomial_chaos
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		p_refinement	Automatic polynomial order refinement
	Optional(*Choose One*)	**Group 1**	askey	Select the Askey basis polynomials that match the basis random variable
			wiener	Use Hermite polynomial basis functions
	Required(*Choose One*)	**Group 2**	quadrature_order	Cubature using tensor-products of Gaussian quadrature rules
			sparse_grid_level	Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids
			cubature_integrand	Cubature using Stroud rules and their extensions
			expansion_order	The maximum order of a polynomial expansion
			orthogonal_least_-interpolation	Build a polynomial chaos expansion from simulation samples using orthogonal least interpolation.

			import_expansion_-file	Build a Polynomial Chaos Expansion (PCE) by import coefficients and a multi-index from a file
	Optional		variance_based_-decomp	Global sensitivity analyis using variance-based decomposition
	Optional*(Choose One)*	**Group 3**	diagonal_-covariance	Display only the diagonal terms of the covariance matrix
			full_covariance	Display the full covariance matrix
	Optional		normalized	The normalized specification requests output of PCE coefficients that correspond to normalized orthogonal basis polynomials
	Optional		sample_type	Choose the sampling algorithm
	Optional		probability_-refinement	Allow refinement of probability and generalized reliability results using importance sampling
	Optional		export_points_file	Specify file to write out the responses from the surrogate
	Optional		export_expansion_-file	Export the coefficients and multi-index of a Polynomial Chaos Expansion (PCE) to a file

	Optional	fixed_seed	Reuse the same seed value for multiple sampling sets
	Optional	reliability_levels	Specify reliability levels at which the response values will be estimated
	Optional	response_levels	Specify the response values at which to estimate relevant statistics
	Optional	distribution	Select between cumulative or complementary-cumulative functions
	Optional	probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional	gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional	rng	Choose a random number generator
	Optional	samples	Specify the number of samples
	Optional	seed	Specify the seed of the random number generator
	Optional	model_pointer	Specify which model block will be used by a method

Description

The polynomial chaos expansion (PCE) is a general framework for the approximate representation of random response functions in terms of finite-dimensional series expansions in standardized random variables. The following provides details on the various polynomial chaos method options in Dakota.

Groups 1 and 2, plus the optional keywords `p_refinement` and `fixed_seed` relate to the specification of a PCE method. In addition, this method treats variables that are not aleatoric-uncertain different, despite the active keyword.

Group 3, and the remainder of the optional keywords relate to the output of the method.

polynomial_chaos Group 1

Group 1 keywords are used to select the type of basis, Ψ_i, of the expansion. Three approaches may be employed:

- Extended (default if no option is selected)
- Wiener
- Askey

The selection of Wiener versus Askey versus Extended is partially automated and partially under the user's control.

- The Extended option is the default and supports only Gaussian correlations.
- If needed to support prescribed correlations (not under user control), the Extended and Askey options will fall back to the Wiener option *on a per variable basis*. If the prescribed correlations are also unsupported by Wiener expansions, then Dakota will exit with an error.
- The Extended option avoids the use of any nonlinear variable transformations by augmenting the Askey approach with numerically-generated orthogonal polynomials for non-Askey probability density functions.
- Extended polynomial selections replace each of the sub-optimal Askey basis selections with numerically-generated polynomials that are orthogonal to the prescribed probability density functions (for bounded normal, lognormal, bounded lognormal, loguniform, triangular, gumbel, frechet, weibull, and bin-based histogram).

polynomial_chaos Group 2
To obtain the coefficients α_i of the expansion, six options are provided:

- multidimensional integration by a tensor-product of Gaussian quadrature rules (specified with `quadrature_order`, and, optionally, `dimension_preference`).
- multidimensional integration by the Smolyak sparse grid method (specified with `sparse_grid_level` and, optionally, `dimension_preference`)
- multidimensional integration by Stroud cubature rules and extensions as specified with `cubature_integrand`.
- multidimensional integration by Latin hypercube sampling (specified with `expansion_samples`).
- linear regression (specified with either `collocation_points` or `collocation_ratio`).
- coefficient import from a file (specified with `import_expansion_file`). A total-order expansion is assumed and must be specified using `expansion_order`.

Active Variables
The default behavior is to form expansions over aleatory uncertain continuous variables. To form expansions over a broader set of variables, one needs to specify `active` followed by `state`, `epistemic`, `design`, or `all` in the variables specification block.

For continuous design, continuous state, and continuous epistemic uncertain variables included in the expansion, Legendre chaos bases are used to model the bounded intervals for these variables. However, these variables are not assumed to have any particular probability distribution, only that they are independent variables. Moreover, when probability integrals are evaluated, only the aleatory random variable domain is integrated, leaving behind a polynomial relationship between the statistics and the remaining design/state/epistemic variables.

polynomial_chaos Group 3
These two keywords are used to specify how this method outputs the covariance of the responses.

Optional Keywords regarding method outputs

Each of these sampling specifications refer to sampling on the PCE approximation for the purposes of generating approximate statistics.

- `sample_type`
- `samples`
- `seed`
- `fixed_seed`
- `rng`
- `probability_refinement`
- `reliability_levels`
- `response_levels`
- `probability_levels`
- `gen_reliability_levels`

which should be distinguished from simulation sampling for generating the PCE coefficients as described in options 4 and 5 above (although options 4 and 5 will share the `sample_type`, `seed`, and `rng` settings, if provided).

When using the `probability_refinement` control, the number of refinement samples is not under the user's control (these evaluations are approximation-based, so management of this expense is less critical). This option allows for refinement of probability and generalized reliability results using importance sampling.

Multi-fidelity UQ

The advanced use case of multifidelity UQ automatically becomes active if the model selected for iteration by the method specification is a multifidelity surrogate model (see hierarchical). In this case, an expansion will first be formed for the model discrepancy (the difference between response results if `additive correction` or the ratio of results if `multiplicative correction`), using the first `quadrature_order` or `sparse_grid_level` value along with any specified refinement strategy. Second, an expansion will be formed for the low fidelity surrogate model, using the second `quadrature_order` or `sparse_grid_level` value (if present; the first is reused if not present) along with any specified refinement strategy. Then the two expansions are combined (added or multiplied) into an expansion that approximates the high fidelity model, from which the final set of statistics are generated. For polynomial chaos expansions, this high fidelity expansion can differ significantly in form from the low fidelity and discrepancy expansions, particularly in the `multiplicative` case where it is expanded to include all of the basis products.

Theory

The polynomial chaos expansion (PCE) is a general framework for the approximate representation of random response functions in terms of finite-dimensional series expansions in standardized random variables

$$R = \sum_{i=0}^{P} \alpha_i \Psi_i(\xi)$$

where α_i is a deterministic coefficient, Ψ_i is a multidimensional orthogonal polynomial and ξ is a vector of standardized random variables. An important distinguishing feature of the methodology is that the functional

relationship between random inputs and outputs is captured, not merely the output statistics as in the case of many nondeterministic methodologies. Dakota provides access to PCE methods through the NonDPolynomialChaos class. Refer to the Uncertainty Quantification Capabilities chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information on the PCE algorithm.

If n is small (e.g., two or three), then tensor-product Gaussian quadrature is quite effective and can be the preferred choice. For moderate to large n (e.g., five or more), tensor-product quadrature quickly becomes too expensive and the sparse grid and regression approaches are preferred. Random sampling for coefficient estimation is generally not recommended due to its slow convergence rate. For incremental studies, approaches 4 and 5 support reuse of previous samples through the incremental_lhs and reuse_points specifications, respectively.

In the quadrature and sparse grid cases, growth rates for nested and non-nested rules can be synchronized for consistency. For a non-nested Gauss rule used within a sparse grid, linear one-dimensional growth rules of $m = 2l + 1$ are used to enforce odd quadrature orders, where l is the grid level and m is the number of points in the rule. The precision of this Gauss rule is then $i = 2m - 1 = 4l + 1$. For nested rules, order growth with level is typically exponential; however, the default behavior is to restrict the number of points to be the lowest order rule that is available that meets the one-dimensional precision requirement implied by either a level l for a sparse grid ($i = 4l + 1$) or an order m for a tensor grid ($i = 2m - 1$). This behavior is known as "restricted growth" or "delayed sequences." To override this default behavior in the case of sparse grids, the `unrestricted` keyword can be used; it cannot be overridden for tensor grids using nested rules since it also provides a mapping to the available nested rule quadrature orders. An exception to the default usage of restricted growth is the `dimension_adaptive` `p_refinement` `generalized` sparse grid case described previously, since the ability to evolve the index sets of a sparse grid in an unstructured manner eliminates the motivation for restricting the exponential growth of nested rules.

See Also

These keywords may also be of interest:

- adaptive_sampling
- gpais
- local_reliability
- global_reliability
- sampling
- importance_sampling
- stoch_collocation

p_refinement

- Keywords Area
- method
- polynomial_chaos
- p_refinement

Automatic polynomial order refinement

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	uniform	Refine an expansion uniformly in all dimensions.
			dimension_-adaptive	Perform automated expansion refinement by adapting to dimensions with higher 'importance'.

Description

The `p_refinement` keyword specifies the usage of automated polynomial order refinement, which can be either `uniform` or `dimension_adaptive`.

The `dimension_adaptive` option is supported for the tensor-product quadrature and Smolyak sparse grid options and `uniform` is supported for tensor and sparse grids as well as regression approaches (`collocation_points` or `collocation_ratio`).

Each of these refinement cases makes use of the `max_iterations` and `convergence_tolerance` method independent controls. The former control limits the number of refinement iterations, and the latter control terminates refinement when the two-norm of the change in the response covariance matrix (or, in goal-oriented approaches, the two-norm of change in the statistical quantities of interest (QOI)) falls below the tolerance.

The `dimension_adaptive` case can be further specified to utilize `sobol`, `decay`, or `generalized` refinement controls. The former two cases employ anisotropic tensor/sparse grids in which the anisotropic dimension preference (leading to anisotropic integrations/expansions with differing refinement levels for different random dimensions) is determined using either total Sobol' indices from variance-based decomposition (`sobol` case: high indices result in high dimension preference) or using spectral coefficient decay rates from a rate estimation technique similar to Richardson extrapolation (`decay` case: low decay rates result in high dimension preference). In these two cases as well as the `uniform` refinement case, the `quadrature_order` or `sparse_grid_level` are ramped by one on each refinement iteration until either of the two convergence controls is satisfied. For the `uniform` refinement case with regression approaches, the `expansion_order` is ramped by one on each iteration while the oversampling ratio (either defined by `collocation_ratio` or inferred from `collocation_points` based on the initial expansion) is held fixed. Finally, the `generalized` `dimension_adaptive` case is the default adaptive approach; it refers to the generalized sparse grid algorithm, a greedy approach in which candidate index sets are evaluated for their impact on the statistical QOI, the most influential sets are selected and used to generate additional candidates, and the index set frontier of a sparse grid is evolved in an unstructured and goal-oriented manner (refer to User's Manual PCE descriptions for additional specifics).

For the case of p_refinement or the case of an explicit nested override, Gauss-Hermite rules are replaced with Genz-Keister nested rules and Gauss-Legendre rules are replaced with Gauss-Patterson nested rules, both of which exchange lower integrand precision for greater point reuse.

uniform

- Keywords Area
- method

- polynomial_chaos
- p_refinement
- uniform

Refine an expansion uniformly in all dimensions.

Specification

Alias: none
Argument(s): none

Description

The quadrature_order or sparse_grid_level are ramped by one on each refinement iteration until either of the two convergence controls is satisfied. For the uniform refinement case with regression approaches, the expansion-_order is ramped by one on each iteration while the oversampling ratio (either defined by collocation_ratio or inferred from collocation_points based on the initial expansion) is held fixed.

dimension_adaptive

- Keywords Area
- method
- polynomial_chaos
- p_refinement
- dimension_adaptive

Perform automated expansion refinement by adapting to dimensions with higher 'importance'.

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	sobol	Determine dimension preference for refinement of a polynomial chaos expansion from the sobol sensitivity indices of the PCE.

			decay	Use spectral coefficient decay rates from a rate estimation technique similar to Richardson extrapolation to control dimension refinement.
			generalized	Use the generalized sparse grid dimension adaptive algorithm to refine a sparse grid approximation of polynomial expansion.

Description

Perform automated expansion refinement by adapting to dimensions with higher 'importance'. Dimension_-adaptive refinement can be further specified as either sobol or generalized (decay not supported). Each of these automated refinement approaches makes use of the max_iterations and convergence_tolerance iteration controls

sobol

- Keywords Area
- method
- polynomial_chaos
- p_refinement
- dimension_adaptive
- sobol

Determine dimension preference for refinement of a polynomial chaos expansion from the sobol sensitivity indices of the PCE.

Specification

Alias: none
Argument(s): none

Description

Determine dimension preference for refinement of a polynomial chaos expansion from the sobol sensitivity indices of the PCE. High indices result in high dimension preference.

decay

- Keywords Area
- method
- polynomial_chaos
- p_refinement
- dimension_adaptive
- decay

Use spectral coefficient decay rates from a rate estimation technique similar to Richardson extrapolation to control dimension refinement.

Specification

Alias: none
Argument(s): none

Description

Use spectral coefficient decay rates from a rate estimation technique similar to Richardson extrapolation to control dimension refinement. Low decay rates result in high dimension preference.

generalized

- Keywords Area
- method
- polynomial_chaos
- p_refinement
- dimension_adaptive
- generalized

Use the generalized sparse grid dimension adaptive algorithm to refine a sparse grid approximation of polynomial expansion.

Specification

Alias: none
Argument(s): none

Description

The generalized sparse grid algorithm is a greedy approach in which candidate index sets are evaluated for their impact on the statistical QOI, the most influential sets are selected and used to generate additional candidates, and the index set frontier of a sparse grid is evolved in an unstructured and goal-oriented manner (refer to User's Manual PCE descriptions for additional specifics).

askey

- Keywords Area
- method
- polynomial_chaos
- askey

Select the Askey basis polynomials that match the basis random variable

Specification

Alias: none
Argument(s): none

Description

The Askey option employs an extended basis of Hermite, Legendre, Laguerre, Jacobi, and generalized Laguerre orthogonal polynomials.

- Askey polynomial selections include Hermite for normal (optimal) as well as bounded normal, lognormal, bounded lognormal, gumbel, frechet, and weibull (sub-optimal); Legendre for uniform (optimal) as well as loguniform, triangular, and bin-based histogram (sub-optimal); Laguerre for exponential (optimal); Jacobi for beta (optimal); and generalized Laguerre for gamma (optimal).

See Also

These keywords may also be of interest:

- wiener

wiener

- Keywords Area
- method
- polynomial_chaos
- wiener

Use Hermite polynomial basis functions

Specification

Alias: none
Argument(s): none

Description

The Wiener option uses a Hermite orthogonal polynomial basis for all random variables and employs the same nonlinear variable transformation as the local and global reliability methods (and therefore has the same variable support).

See Also

These keywords may also be of interest:

- askey

quadrature_order

- Keywords Area
- method
- polynomial_chaos
- quadrature_order

Cubature using tensor-products of Gaussian quadrature rules

Specification

Alias: none

Argument(s): INTEGERLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		dimension_-preference	A set of weights specifying the realtive importance of each uncertain variable (dimension)
	Optional*(Choose One)*	**Group 1**	nested	Enforce use of nested quadrature rules if available
			non_nested	Enforce use of non-nested quadrature rules

Description

Multidimensional integration by a tensor-product of Gaussian quadrature rules (specified with `quadrature_order`, and, optionally, `dimension_preference`). The default rule selection is to employ `non_nested` Gauss rules including Gauss-Hermite (for normals or transformed normals), Gauss-Legendre (for uniforms or transformed uniforms), Gauss-Jacobi (for betas), Gauss-Laguerre (for exponentials), generalized Gauss-Laguerre (for gammas), and numerically-generated Gauss rules (for other distributions when using an Extended basis). For the case of `p_refinement` or the case of an explicit `nested` override, Gauss-Hermite rules are replaced with Genz-Keister nested rules and Gauss-Legendre rules are replaced with Gauss-Patterson nested rules, both of which exchange lower integrand precision for greater point reuse. By specifying a `dimension_preference`, where higher preference leads to higher order polynomial resolution, the tensor grid may be rendered anisotropic. The dimension specified to have highest preference will be set to the specified `quadrature_order` and all other dimensions will be reduced in proportion to their reduced preference; any non-integral portion is truncated. To synchronize with tensor-product integration, a tensor-product expansion is used, where the order p_i of the expansion in each dimension is selected to be half of the integrand precision available from the rule in use,

rounded down. In the case of non-nested Gauss rules with integrand precision $2m_i - 1$, p_i is one less than the quadrature order m_i in each dimension (a one-dimensional expansion contains the same number of terms, $p + 1$, as the number of Gauss points). The total number of terms, *N*, in a tensor-product expansion involving *n* uncertain input variables is

$$N = 1 + P = \prod_{i=1}^{n} (p_i + 1)$$

In some advanced use cases (e.g., multifidelity UQ), multiple grid resolutions can be employed; for this reason, the `quadrature_order` specification supports an array input.

dimension_preference

- Keywords Area
- method
- polynomial_chaos
- quadrature_order
- dimension_preference

A set of weights specifying the realtive importance of each uncertain variable (dimension)

Specification

Alias: none
Argument(s): REALLIST

Description

A set of weights specifying the realtive importance of each uncertain variable (dimension). Using this specification leada to anisotropic integrations with differing refinement levels for different random dimensions.

See Also

These keywords may also be of interest:

- sobol
- decay

nested

- Keywords Area
- method
- polynomial_chaos
- quadrature_order
- nested

Enforce use of nested quadrature rules if available

Specification

Alias: none
Argument(s): none

Description

Enforce use of nested quadrature rules if available. For instance if the aleatory variables are Gaussian use the Nested Genz-Keister rule instead of the default non-nested Gauss-Hermite rule variables are

non_nested

- Keywords Area
- method
- polynomial_chaos
- quadrature_order
- non_nested

Enforce use of non-nested quadrature rules

Specification

Alias: none
Argument(s): none

Description

Enforce use of non-nested quadrature rules if available. For instance if the aleatory variables are Gaussian use the non-nested Gauss-Hermite rule

sparse_grid_level

- Keywords Area
- method
- polynomial_chaos
- sparse_grid_level

Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Specification

Alias: none
Argument(s): INTEGERLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	restricted	Restrict the growth rates for nested and non-nested rules can be synchronized for consistency.
			unrestricted	Overide the default restriction of growth rates for nested and non-nested rules that are by defualt synchronized for consistency.
	Optional		dimension_-preference	A set of weights specifying the realtive importance of each uncertain variable (dimension)
	Optional*(Choose One)*	**Group 2**	nested	Enforce use of nested quadrature rules if available
			non_nested	Enforce use of non-nested quadrature rules

Description

Multi-dimensional integration by the Smolyak sparse grid method (specified with sparse_grid_level and, optionally, dimension_preference). The underlying one-dimensional integration rules are the same as for the tensor-product quadrature case; however, the default rule selection is nested for sparse grids (Genz-Keister for normals/transformed normals and Gauss-Patterson for uniforms/transformed uniforms). This default can be overridden with an explicit non_nested specification (resulting in Gauss-Hermite for normals/transformed normals and Gauss-Legendre for uniforms/transformed uniforms). As for tensor quadrature, the dimension_preference specification enables the use of anisotropic sparse grids (refer to the PCE description in the User's Manual for the anisotropic index set constraint definition). Similar to anisotropic tensor grids, the dimension with greatest preference will have resolution at the full sparse_grid_level and all other dimension resolutions will be reduced in proportion to their reduced preference. For PCE with either isotropic or anisotropic sparse grids, a summation of tensor-product expansions is used, where each anisotropic tensor-product quadrature rule underlying the sparse grid construction results in its own anisotropic tensor-product expansion as described in case 1. These anisotropic tensor-product expansions are summed into a sparse PCE using the standard Smolyak summation (again, refer to the User's Manual for additional details). As for quadrature_order, the sparse_grid_level specification admits an array input for enabling specification of multiple grid resolutions used by certain advanced solution methodologies.

This keyword can be used when using sparse grid integration to calculate PCE coefficients or when generating samples for sparse grid collocation.

restricted

- Keywords Area
- method
- polynomial_chaos
- sparse_grid_level
- restricted

Restrict the growth rates for nested and non-nested rules can be synchronized for consistency.

Specification

Alias: none
Argument(s): none

Description

In the quadrature and sparse grid cases, growth rates for nested and non-nested rules can be synchronized for consistency. For a non-nested Gauss rule used within a sparse grid, linear one-dimensional growth rules of $m = 2l + 1$ are used to enforce odd quadrature orders, where l is the grid level and m is the number of points in the rule. The precision of this Gauss rule is then $i = 2m - 1 = 4l + 1$. For nested rules, order growth with level is typically exponential; however, the default behavior is to restrict the number of points to be the lowest order rule that is available that meets the one-dimensional precision requirement implied by either a level l for a sparse grid ($i = 4l + 1$) or an order m for a tensor grid ($i = 2m - 1$). This behavior is known as "restricted growth" or "delayed sequences." To override this default behavior in the case of sparse grids, the unrestricted keyword can be used; it cannot be overridden for tensor grids using nested rules since it also provides a mapping to the available nested rule quadrature orders. An exception to the default usage of restricted growth is the dimension_adaptive p-_refinement generalized sparse grid case described previously, since the ability to evolve the index sets of a sparse grid in an unstructured manner eliminates the motivation for restricting the exponential growth of nested rules.

unrestricted

- Keywords Area
- method
- polynomial_chaos
- sparse_grid_level
- unrestricted

Overide the default restriction of growth rates for nested and non-nested rules that are by defualt synchronized for consistency.

Specification

Alias: none
Argument(s): none

Description

In the quadrature and sparse grid cases, growth rates for nested and non-nested rules can be synchronized for consistency. For a non-nested Gauss rule used within a sparse grid, linear one-dimensional growth rules of $m = 2l + 1$ are used to enforce odd quadrature orders, where l is the grid level and m is the number of points in the rule. The precision of this Gauss rule is then $i = 2m - 1 = 4l + 1$. For nested rules, order growth with level is typically exponential; however, the default behavior is to restrict the number of points to be the lowest order rule that is available that meets the one-dimensional precision requirement implied by either a level l for a sparse grid ($i = 4l + 1$) or an order m for a tensor grid ($i = 2m - 1$). This behavior is known as "restricted growth" or "delayed sequences." To override this default behavior in the case of sparse grids, the unrestricted keyword can be used; it cannot be overridden for tensor grids using nested rules since it also provides a mapping to the available nested rule quadrature orders. An exception to the default usage of restricted growth is the dimension_adaptive p-_refinement generalized sparse grid case described previously, since the ability to evolve the index sets of a sparse grid in an unstructured manner eliminates the motivation for restricting the exponential growth of nested rules.

dimension_preference

- Keywords Area
- method
- polynomial_chaos
- sparse_grid_level
- dimension_preference

A set of weights specifying the realtive importance of each uncertain variable (dimension)

Specification

Alias: none
Argument(s): REALLIST

Description

A set of weights specifying the realtive importance of each uncertain variable (dimension). Using this specification leada to anisotropic integrations with differing refinement levels for different random dimensions.

See Also

These keywords may also be of interest:

- sobol
- decay

nested

- Keywords Area
- method
- polynomial_chaos
- sparse_grid_level
- nested

Enforce use of nested quadrature rules if available

Specification

Alias: none
Argument(s): none

Description

Enforce use of nested quadrature rules if available. For instance if the aleatory variables are Gaussian use the Nested Genz-Keister rule instead of the default non-nested Gauss-Hermite rule variables are

non_nested

- Keywords Area
- method
- polynomial_chaos
- sparse_grid_level
- non_nested

Enforce use of non-nested quadrature rules

Specification

Alias: none
Argument(s): none

Description

Enforce use of non-nested quadrature rules if available. For instance if the aleatory variables are Gaussian use the non-nested Gauss-Hermite rule

cubature_integrand

- Keywords Area
- method
- polynomial_chaos
- cubature_integrand

Cubature using Stroud rules and their extensions

Specification

Alias: none
Argument(s): INTEGER

Description

Multi-dimensional integration by Stroud cubature rules [[78] "Stroud, 1971"] and extensions [[91] "Xiu, 2008"], as specified with `cubature_integrand`. A total-order expansion is used, where the isotropic order p of the expansion is half of the integrand order, rounded down. The total number of terms N for an isotropic total-order expansion of order p over n variables is given by

$$N = 1 + P = 1 + \sum_{s=1}^{p} \frac{1}{s!} \prod_{r=0}^{s-1} (n+r) = \frac{(n+p)!}{n!p!}$$

Since the maximum integrand order is currently five for normal and uniform and two for all other types, at most second- and first-order expansions, respectively, will be used. As a result, cubature is primarily useful for global sensitivity analysis, where the Sobol' indices will provide main effects and, at most, two-way interactions. In addition, the random variable set must be independent and identically distributed (*iid*), so the use of `askey` or `wiener` transformations may be required to create *iid* variable sets in the transformed space (as well as to allow usage of the higher order cubature rules for normal and uniform). Note that global sensitivity analysis often assumes uniform bounded regions, rather than precise probability distributions, so the *iid* restriction would not be problematic in that case.

expansion_order

- Keywords Area
- method
- polynomial_chaos
- expansion_order

The maximum order of a polynomial expansion

Specification

Alias: none
Argument(s): INTEGERLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		dimension_-preference	A set of weights specifying the realtive importance of each uncertain variable (dimension)

	Optional		basis_type	Specify the type of truncation to be used with a Polynomial Chaos Expansion.
	Required*(Choose One)*	**Group 1**	collocation_points	Specify the number of collocation points used to estimate PCE coefficients using regression or orthogonal-least-interpolation.
			collocation_ratio	Set the number of points used to build a PCE via regression to be proportional to the number of terms in the expansion.
			expansion_samples	The Number simulation samples to estimate the PCE coefficients
	Optional		import_points_file	Specify file containing points (variable and response values)

Description

Multidimensional integration by Latin hypercube sampling (specified with `expansion_samples`). In this case, the expansion order p cannot be inferred from the numerical integration specification and it is necessary to provide an `expansion_order` to specify p for a total-order expansion.

Linear regression (specified with either `collocation_points` or `collocation_ratio`). A total-order expansion is used and must be specified using `expansion_order` as described in the previous option. To avoid requiring the user to calculate N from n and p), the `collocation_ratio` allows for specification of a constant factor applied to N (e.g., `collocation_ratio` = `2`. produces samples = $2N$). In addition, the default linear relationship with N can be overridden using a real-valued exponent specified using `ratio_order`. In this case, the number of samples becomes cN^o where c is the `collocation_ratio` and o is the `ratio_order`. The `use_derivatives` flag informs the regression approach to include derivative matching equations (limited to gradients at present) in the least squares solutions, enabling the use of fewer collocation points for a given expansion order and dimension (number of points required becomes $\frac{cN^o}{n+1}$). When admissible, a constrained least squares approach is employed in which response values are first reproduced exactly and error in reproducing response derivatives is minimized. Two collocation grid options are supported: the default is Latin hypercube sampling (”point collocation”), and an alternate approach of ”probabilistic collocation” is also available through inclusion of the `tensor_grid` keyword. In this alternate case, the collocation grid is defined using a subset of tensor-product quadrature points: the order of the tensor-product grid is selected as one more than the expansion order in each dimension (to avoid sampling at roots of the basis polynomials) and then the tensor multi-index is uniformly sampled to generate a non-repeated subset of tensor quadrature points.

If `collocation_points` or `collocation_ratio` is specified, the PCE coefficients will be determined by regression. If no regression specification is provided, appropriate defaults are defined. Specifically SVD-based least-squares will be used for solving over-determined systems and under-determined systems will be solved using LASSO. For the situation when the number of function values is smaller than the number of terms in a PCE, but the total number of samples including gradient values is greater than the number of terms, the resulting over-determined system will be solved using equality constrained least squares. Technical information on the various methods listed below can be found in the Linear regression section of the Theory Manual. Some of the regression methods (OMP, LASSO, and LARS) are able to produce a set of possible PCE coefficient vectors (see the Linear regression section in the Theory Manual). If cross validation is inactive, then only one solution, consistent with the `noise_tolerance`, will be returned. If cross validation is active, Dakota will choose between possible coefficient vectors found internally by the regression method across the set of expansion orders (1,...,`expansion_order`) and the set of specified noise tolerances and return the one with the lowest cross validation error indicator.

dimension_preference

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- dimension_preference

A set of weights specifying the realtive importance of each uncertain variable (dimension)

Specification

Alias: none
Argument(s): REALLIST

Description

A set of weights specifying the realtive importance of each uncertain variable (dimension). Using this specification leada to anisotropic integrations with differing refinement levels for different random dimensions.

See Also

These keywords may also be of interest:

- sobol
- decay

basis_type

- Keywords Area
- method
- polynomial_chaos

- expansion_order
- basis_type

Specify the type of truncation to be used with a Polynomial Chaos Expansion.

Specification

Alias: none
Argument(s): none

	Required/-Optional	**Description of Group**	**Dakota Keyword**	**Dakota Keyword Description**
	Required*(Choose One)*	**Group 1**	tensor_product	Use a tensor-product index set to construct a polynomial chaos expansion.
			total_order	Use a total-order index set to construct a polynomial chaos expansion.
			adapted	Use adaptive basis selection to choose the basis terms in a polynomial chaos expansion.

Description

Specify the type of truncation to be used with a Polynomial Chaos Expansion.

tensor_product

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- basis_type
- tensor_product

Use a tensor-product index set to construct a polynomial chaos expansion.

Specification

Alias: none
Argument(s): none

Description

Use a tensor-product index set to construct a polynomial chaos expansion. That is for a given order p keep all terms with d-dimensional multi index $\mathbf{i} = (i_1, \ldots, i_d)$ that satisfies

$$\max(i_1, \ldots, i_d) \leq p$$

total_order

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- basis_type
- total_order

Use a total-order index set to construct a polynomial chaos expansion.

Specification

Alias: none
Argument(s): none

Description

Use the traditional total-order index set to construct a polynomial chaos expansion. That is for a given order p keep all terms with a d-dimensional multi index $\mathbf{i} = (i_1, \ldots, i_d)$ that satisfies

$$\sum_{k=1}^{d} i_k \leq p$$

adapted

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- basis_type
- adapted

Use adaptive basis selection to choose the basis terms in a polynomial chaos expansion.

Specification

Alias: none
Argument(s): none

Description

Use adaptive basis selection to choose the basis terms in a polynomial chaos expansion.

collocation_points

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points

Specify the number of collocation points used to estimate PCE coefficients using regression or orthogonal-least-interpolation.

Specification

Alias: none
Argument(s): INTEGERLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		ratio_order	Specify a non-linear the relationship between the expansion order of a polynomial chaos expansion and the number of samples that will be used to compute the PCE coefficients.

	Optional*(Choose One)*	**Group 1**	least_squares	Compute the coefficients of a polynomial expansion using least squares
			orthogonal_-matching_pursuit	Compute the coefficients of a polynomial expansion using orthogonal matching pursuit (OMP)
			basis_pursuit	Compute the coefficients of a polynomial expansion by solving the Basis Pursuit ℓ_1-minimization problem using linear programming.
			basis_pursuit_-denoising	Compute the coefficients of a polynomial expansion by solving the Basis Pursuit Denoising ℓ_1-minimization problem using second order cone optimization.
			least_angle_-regression	Compute the coefficients of a polynomial expansion by using the greedy least angle regression (LAR) method.

			least_absolute_-shrinkage	Compute the coefficients of a polynomial expansion by using the LASSO problem.
	Optional		cross_validation	Use cross validation to choose the 'best' polynomial order of a polynomial chaos expansion.
	Optional		use_derivatives	Use derivative data to construct surrogate models
	Optional		tensor_grid	Use sub-sampled tensor-product quadrature points to build a polynomial chaos expansion.
	Optional		reuse_points	??

Description

Specify the number of collocation points used to estimate PCE coefficients using regression or orthogonal-least-interpolation.

ratio_order

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- ratio_order

Specify a non-linear the relationship between the expansion order of a polynomial chaos expansion and the number of samples that will be used to compute the PCE coefficients.

Specification

Alias: none
Argument(s): REAL

Description

When using regression type methods (specified with either collocation_points or collocation_ratio), a total-order expansion can be specified using expansion_order. To avoid requiring the user to calculate N from n and p), the collocation_ratio allows for specification of a constant factor applied to N (e.g., collocation_ratio = 2. produces samples = 2N). In addition, the default linear relationship with N can be overridden using a real-valued exponent specified using ratio_order. In this case, the number of samples becomes cN^o where c is the collocation_-ratio and o is the ratio_order.

least_squares

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- least_squares

Compute the coefficients of a polynomial expansion using least squares

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	svd	Calculate the coefficients of a polynomial chaos expansion using the singular value decomposition.
			equality_-constrained	Calculate the coefficients of a polynomial chaos expansion using equality constrained least squares.

Description

Compute the coefficients of a polynomial expansion using least squares. Specifically SVD-based least-squares will be used for solving over-determined systems. For the situation when the number of function values is smaller than the number of terms in a PCE, but the total number of samples including gradient values is greater than the number of terms, the resulting over-determined system will be solved using equality constrained least squares

svd

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- least_squares
- svd

Calculate the coefficients of a polynomial chaos expansion using the singular value decomposition.

Specification

Alias: none
Argument(s): none

Description

Calculate the coefficients of a polynomial chaos expansion using the singular value decomposition. When the number of model runs exceeds the number of terms in the PCE, the solution returned will be the least-squares solution, otherwise the solution will be the minimum norm solution computed using the pseudo-inverse.

equality_constrained

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- least_squares
- equality_constrained

Calculate the coefficients of a polynomial chaos expansion using equality constrained least squares.

Specification

Alias: none
Argument(s): none

Description

Calculate the coefficients of a polynomial chaos expansion using equality constrained least squares.

orthogonal_matching_pursuit

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- orthogonal_matching_pursuit

Compute the coefficients of a polynomial expansion using orthogonal matching pursuit (OMP)

Specification

Alias: omp

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		noise_tolerance	The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Description

Compute the coefficients of a polynomial expansion using orthogonal matching pursuit (OMP). Orthogonal matching pursuit (OMP) is a greedy algorithm that is usefult when solving underdetermined linear systems.

noise_tolerance

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- orthogonal_matching_pursuit
- noise_tolerance

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Specification

Alias: none
Argument(s): REALLIST

Description

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

basis_pursuit

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- basis_pursuit

Compute the coefficients of a polynomial expansion by solving the Basis Pursuit ℓ_1-minimization problem using linear programming.

Specification

Alias: bp
Argument(s): none

Description

Compute the coefficients of a polynomial expansion by solving the Basis Pursuit ℓ_1-minimization problem using linear programming.

basis_pursuit_denoising

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- basis_pursuit_denoising

Compute the coefficients of a polynomial expansion by solving the Basis Pursuit Denoising ℓ_1-minimization problem using second order cone optimization.

Specification

Alias: bpdn
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		noise_tolerance	The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Description

Compute the coefficients of a polynomial expansion by solving the Basis Pursuit Denoising ℓ_1-minimization problem using second order cone optimization.

noise_tolerance

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- basis_pursuit_denoising
- noise_tolerance

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Specification

Alias: none
Argument(s): REALLIST

Description

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

least_angle_regression

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- least_angle_regression

Compute the coefficients of a polynomial expansion by using the greedy least angle regression (LAR) method.

Specification

Alias: lars
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		noise_tolerance	The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Description

Compute the coefficients of a polynomial expansion by using the greedy least angle regression (LAR) method.

noise_tolerance

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- least_angle_regression
- noise_tolerance

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Specification

Alias: none
Argument(s): REALLIST

Description

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

least_absolute_shrinkage

- Keywords Area
- method
- polynomial_chaos
- expansion_order

- collocation_points
- least_absolute_shrinkage

Compute the coefficients of a polynomial expansion by using the LASSO problem.

Specification

Alias: lasso
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		noise_tolerance	The noise tolerance used when performing cross validation in the presence of noise or truncation errors.
	Optional		l2_penalty	The l_2 pentalty used when performing compressed sensing with elastic net.

Description

Compute the coefficients of a polynomial expansion by using the LASSO problem.

noise_tolerance

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- least_absolute_shrinkage
- noise_tolerance

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Specification

Alias: none
Argument(s): REALLIST

Description

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

l2_penalty

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- least_absolute_shrinkage
- l2_penalty

The l_2 pentalty used when performing compressed sensing with elastic net.

Specification

Alias: none
Argument(s): REAL

Description

The l_2 pentalty used when performing compressed sensing with elastic net.

cross_validation

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- cross_validation

Use cross validation to choose the 'best' polynomial order of a polynomial chaos expansion.

Specification

Alias: none
Argument(s): none

Description

Use cross validation to choose the 'best' polynomial degree of a polynomial chaos expansion. 10 fold cross validation is used to estimate the cross validation error of a total-order polynomial expansion for orders 1 through to order. The order chosen is the one that produces the lowest cross validation error. If there are not enough points to perform 10 fold cross validation then one-at-a-time cross validation will be performed.

use_derivatives

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none
Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

tensor_grid

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- tensor_grid

Use sub-sampled tensor-product quadrature points to build a polynomial chaos expansion.

Specification

Alias: none
Argument(s): none

Description

Tthe collocation grid is defined using a subset of tensor-product quadrature points: the order of the tensor-product grid is selected as one more than the expansion order in each dimension (to avoid sampling at roots of the basis polynomials) and then the tensor multi-index is uniformly sampled to generate a non-repeated subset of tensor quadrature points.

reuse_points

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_points
- reuse_points

??

Topics

This keyword is related to the topics:

- problem

Specification

Alias: reuse_samples
Argument(s): none

Description

Similar to reuse_points??

collocation_ratio

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio

Set the number of points used to build a PCE via regression to be proportional to the number of terms in the expansion.

Specification

Alias: none
Argument(s): REAL

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		ratio_order	Specify a non-linear the relationship between the expansion order of a polynomial chaos expansion and the number of samples that will be used to compute the PCE coefficients.
	Optional*(Choose One)*	**Group 1**	least_squares	Compute the coefficients of a polynomial expansion using least squares
			orthogonal_-matching_pursuit	Compute the coefficients of a polynomial expansion using orthogonal matching pursuit (OMP)
			basis_pursuit	Compute the coefficients of a polynomial expansion by solving the Basis Pursuit ℓ_1-minimization problem using linear programming.

			basis_pursuit_-denoising	Compute the coefficients of a polynomial expansion by solving the Basis Pursuit Denoising ℓ_1-minimization problem using second order cone optimization.
			least_angle_-regression	Compute the coefficients of a polynomial expansion by using the greedy least angle regression (LAR) method.
			least_absolute_-shrinkage	Compute the coefficients of a polynomial expansion by using the LASSO problem.
	Optional		cross_validation	Use cross validation to choose the 'best' polynomial order of a polynomial chaos expansion.
	Optional		use_derivatives	Use derivative data to construct surrogate models
	Optional		tensor_grid	Use sub-sampled tensor-product quadrature points to build a polynomial chaos expansion.

	Optional	reuse_points	??

Description

Set the number of points used to build a PCE via regression to be proportional to the number of terms in the expansion. To avoid requiring the user to calculate N from n and p, the collocation_ratio allows for specification of a constant factor applied to N (e.g., collocation_ratio = 2. produces samples = 2N). In addition, the default linear relationship with N can be overridden using a real-valued exponent specified using ratio_order. In this case, the number of samples becomes cN^o where c is the collocation_ratio and o is the ratio_order. The use_derivatives flag informs the regression approach to include derivative matching equations (limited to gradients at present) in the least squares solutions, enabling the use of fewer collocation points for a given expansion order and dimension (number of points required becomes $\frac{cN^o}{n+1}$).

ratio_order

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- ratio_order

Specify a non-linear the relationship between the expansion order of a polynomial chaos expansion and the number of samples that will be used to compute the PCE coefficients.

Specification

Alias: none
Argument(s): REAL

Description

When using regression type methods (specified with either collocation_points or collocation_ratio), a total-order expansion can be specified using expansion_order. To avoid requiring the user to calculate N from n and p), the collocation_ratio allows for specification of a constant factor applied to N (e.g., collocation_ratio = 2. produces samples = 2N). In addition, the default linear relationship with N can be overridden using a real-valued exponent specified using ratio_order. In this case, the number of samples becomes cN^o where c is the collocation_ratio and o is the ratio_order.

least_squares

- Keywords Area
- method
- polynomial_chaos
- expansion_order

- collocation_ratio
- least_squares

Compute the coefficients of a polynomial expansion using least squares

Specification

Alias: none

Argument(s): none

	Required/-Optional **Optional***(Choose One)*	Description of Group **Group 1**	Dakota Keyword svd	Dakota Keyword Description Calculate the coefficients of a polynomial chaos expansion using the singular value decomposition.
			equality_-constrained	Calculate the coefficients of a polynomial chaos expansion using equality constrained least squares.

Description

Compute the coefficients of a polynomial expansion using least squares. Specifically SVD-based least-squares will be used for solving over-determined systems. For the situation when the number of function values is smaller than the number of terms in a PCE, but the total number of samples including gradient values is greater than the number of terms, the resulting over-determined system will be solved using equality constrained least squares

svd

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- least_squares
- svd

Calculate the coefficients of a polynomial chaos expansion using the singular value decomposition.

Specification

Alias: none
Argument(s): none

Description

Calculate the coefficients of a polynomial chaos expansion using the singular value decomposition. When the number of model runs exceeds the number of terms in the PCE, the solution returned will be the least-squares solution, otherwise the solution will be the minimum norm solution computed using the pseudo-inverse.

equality_constrained

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- least_squares
- equality_constrained

Calculate the coefficients of a polynomial chaos expansion using equality constrained least squares.

Specification

Alias: none
Argument(s): none

Description

Calculate the coefficients of a polynomial chaos expansion using equality constrained least squares.

orthogonal_matching_pursuit

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- orthogonal_matching_pursuit

Compute the coefficients of a polynomial expansion using orthogonal matching pursuit (OMP)

Specification

Alias: omp
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		noise_tolerance	The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Description

Compute the coefficients of a polynomial expansion using orthogonal matching pursuit (OMP). Orthogonal matching pursuit (OMP) is a greedy algorithm that is usefult when solving underdetermined linear systems.

noise_tolerance

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- orthogonal_matching_pursuit
- noise_tolerance

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Specification

Alias: none
Argument(s): REALLIST

Description

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

basis_pursuit

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- basis_pursuit

Compute the coefficients of a polynomial expansion by solving the Basis Pursuit ℓ_1-minimization problem using linear programming.

Specification

Alias: bp
Argument(s): none

Description

Compute the coefficients of a polynomial expansion by solving the Basis Pursuit ℓ_1-minimization problem using linear programming.

basis_pursuit_denoising

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- basis_pursuit_denoising

Compute the coefficients of a polynomial expansion by solving the Basis Pursuit Denoising ℓ_1-minimization problem using second order cone optimization.

Specification

Alias: bpdn
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		noise_tolerance	The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Description

Compute the coefficients of a polynomial expansion by solving the Basis Pursuit Denoising ℓ_1-minimization problem using second order cone optimization.

noise_tolerance

- Keywords Area
- method
- polynomial_chaos

- expansion_order
- collocation_ratio
- basis_pursuit_denoising
- noise_tolerance

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Specification

Alias: none
Argument(s): REALLIST

Description

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

least_angle_regression

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- least_angle_regression

Compute the coefficients of a polynomial expansion by using the greedy least angle regression (LAR) method.

Specification

Alias: lars
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		noise_tolerance	The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Description

Compute the coefficients of a polynomial expansion by using the greedy least angle regression (LAR) method.

noise_tolerance

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- least_angle_regression
- noise_tolerance

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Specification

Alias: none
Argument(s): REALLIST

Description

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

least_absolute_shrinkage

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- least_absolute_shrinkage

Compute the coefficients of a polynomial expansion by using the LASSO problem.

Specification

Alias: lasso
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		noise_tolerance	The noise tolerance used when performing cross validation in the presence of noise or truncation errors.
	Optional		l2_penalty	The l_2 pentalty used when performing compressed sensing with elastic net.

Description

Compute the coefficients of a polynomial expansion by using the LASSO problem.

noise_tolerance

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- least_absolute_shrinkage
- noise_tolerance

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

Specification

Alias: none
Argument(s): REALLIST

Description

The noise tolerance used when performing cross validation in the presence of noise or truncation errors.

l2_penalty

- Keywords Area
- method
- polynomial_chaos

- expansion_order
- collocation_ratio
- least_absolute_shrinkage
- l2_penalty

The l_2 pentalty used when performing compressed sensing with elastic net.

Specification

Alias: none
Argument(s): REAL

Description

The l_2 pentalty used when performing compressed sensing with elastic net.

cross_validation

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- cross_validation

Use cross validation to choose the 'best' polynomial order of a polynomial chaos expansion.

Specification

Alias: none
Argument(s): none

Description

Use cross validation to choose the 'best' polynomial degree of a polynomial chaos expansion. 10 fold cross validation is used to estimate the cross validation error of a total-order polynomial expansion for orders 1 through to order. The order chosen is the one that produces the lowest cross validation error. If there are not enough points to perform 10 fold cross validation then one-at-a-time cross validation will be performed.

use_derivatives

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none
Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

tensor_grid

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- tensor_grid

Use sub-sampled tensor-product quadrature points to build a polynomial chaos expansion.

Specification

Alias: none
Argument(s): none

Description

Tthe collocation grid is defined using a subset of tensor-product quadrature points: the order of the tensor-product grid is selected as one more than the expansion order in each dimension (to avoid sampling at roots of the basis polynomials) and then the tensor multi-index is uniformly sampled to generate a non-repeated subset of tensor quadrature points.

reuse_points

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- collocation_ratio
- reuse_points

??

Topics

This keyword is related to the topics:

- problem

Specification

Alias: reuse_samples
Argument(s): none

Description

Similar to reuse_points??

expansion_samples

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- expansion_samples

The Number simulation samples to estimate the PCE coefficients

Specification

Alias: none
Argument(s): INTEGERLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Optional	reuse_points	??
	Optional	incremental_lhs	Augment an existing LHS study

Description

The Number simulation samples to estimate the PCE coefficients In this case, the expansion order p cannot be inferred from the numerical integration specification and it is necessary to provide an expansion_order to specify p for a total-order expansion.

reuse_points

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- expansion_samples
- reuse_points

??

Topics

This keyword is related to the topics:

- problem

Specification

Alias: reuse_samples
Argument(s): none

Description

Similar to reuse_points??

incremental_lhs

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- expansion_samples
- incremental_lhs

Augment an existing LHS study

Specification

Alias: none
Argument(s): none

Description

`incremental_lhs` will augment an existing LHS sampling study with more samples to get better estimates of mean, variance, and percentiles.

The number of samples in the second set MUST currently be 2 times the number of previous samples, although incremental sampling based on any power of two may be supported in future releases.

To use this `sample_type`, a previous LHS (or incremental LHS) sampling study with sample size N must have already been performed, and **the dakota restart file must be available from this previous study.** The variables and responses specifications must be the same in both studies.

See the examples below and the Usage and Restarting Dakota Studies pages.

Examples

For example, say a user performs an initial study using `lhs` as the `sample_type`, and generates 10 samples.

One way to ensure the restart file is saved is to specify a non-default name, via a command line option:

```
dakota -i LHS_10.in -w LHS_10.rst
```

which uses the input file:

```
# LHS_10.in

environment
  tabular_graphics_data
    tabular_graphics_file = 'lhs10.dat'

method
  sampling
    sample_type lhs
    samples = 10

model
  single

variables
  uniform_uncertain = 2
    descriptors  =   'input1'     'input2'
    lower_bounds =  -2.0     -2.0
    upper_bounds =   2.0      2.0

interface
  analysis_drivers 'text_book'
    fork

responses
  response_functions = 1
  no_gradients
  no_hessians
```

and the restart file is written to LHS_10.rst.

Then an incremental LHS study can be run with:

```
dakota -i LHS_20.in -r LHS_10.rst -w LHS_20.rst
```

where `LHS_20.in` is shown below, and LHS_10.rst is the restart file containing the results of the previous LHS study.

```
# LHS_20.in

environment
  tabular_graphics_data
    tabular_graphics_file = 'lhs_incremental_20.dat'

method
  sampling
    sample_type incremental_lhs
    samples = 20
    previous_samples = 10

model
  single

variables
  uniform_uncertain = 2
    descriptors  =   'input1'     'input2'
    lower_bounds =  -2.0     -2.0
    upper_bounds =   2.0      2.0

interface
  analysis_drivers 'text_book'
    fork

responses
  response_functions = 1
  no_gradients
  no_hessians
```

The user will get 10 new LHS samples which maintain both the correlation and stratification of the original LHS sample. The new samples will be combined with the original samples to generate a combined sample of size 20.

This is clearly seen by comparing the two tabular graphics data files.

import_points_file

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none

Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- polynomial_chaos
- expansion_order
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.

- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

orthogonal_least_interpolation

- Keywords Area
- method
- polynomial_chaos
- orthogonal_least_interpolation

Build a polynomial chaos expansion from simulation samples using orthogonal least interpolation.

Specification

Alias: least_interpolation oli

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		collocation_points	Specify the number of collocation points used to estimate PCE coefficients using orthogonal least interpolation
	Optional		cross_validation	Use cross validation to choose the 'best' polynomial order of a polynomial chaos expansion.
	Optional		tensor_grid	Use sub-sampled tensor-product quadrature points to build a polynomial chaos expansion.

	Optional	reuse_points	??
	Optional	import_points_file	Specify file containing points (variable and response values)

Description

Build a polynomial chaos expansion from simulation samples using orthogonal least interpolation. Unlike the other regression methods `expansion_order` cannot be set. OLI will produce the lowest degree polynomial that interpolates the data

collocation_points

- Keywords Area
- method
- polynomial_chaos
- orthogonal_least_interpolation
- collocation_points

Specify the number of collocation points used to estimate PCE coefficients using orthogonal least interpolation

Specification

Alias: none
Argument(s): INTEGERLIST

Description

Specify the number of collocation points used to estimate PCE coefficients using orthogonal least interpolation

cross_validation

- Keywords Area
- method
- polynomial_chaos
- orthogonal_least_interpolation
- cross_validation

Use cross validation to choose the 'best' polynomial order of a polynomial chaos expansion.

Specification

Alias: none
Argument(s): none

Description

Use cross validation to choose the 'best' polynomial degree of a polynomial chaos expansion. 10 fold cross validation is used to estimate the cross validation error of a total-order polynomial expansion for orders 1 through to order. The order chosen is the one that produces the lowest cross validation error. If there are not enough points to perform 10 fold cross validation then one-at-a-time cross validation will be performed.

tensor_grid

- Keywords Area
- method
- polynomial_chaos
- orthogonal_least_interpolation
- tensor_grid

Use sub-sampled tensor-product quadrature points to build a polynomial chaos expansion.

Specification

Alias: none
Argument(s): INTEGERLIST

Description

Tthe collocation grid is defined using a subset of tensor-product quadrature points: the order of the tensor-product grid is selected as one more than the expansion order in each dimension (to avoid sampling at roots of the basis polynomials) and then the tensor multi-index is uniformly sampled to generate a non-repeated subset of tensor quadrature points.

reuse_points

- Keywords Area
- method
- polynomial_chaos
- orthogonal_least_interpolation
- reuse_points

??

Topics

This keyword is related to the topics:

- problem

Specification

Alias: reuse_samples
Argument(s): none

Description

Similar to reuse_points??

import_points_file

- Keywords Area
- method
- polynomial_chaos
- orthogonal_least_interpolation
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none
Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- polynomial_chaos
- orthogonal_least_interpolation
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- polynomial_chaos
- orthogonal_least_interpolation
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

import_expansion_file

- Keywords Area
- method
- polynomial_chaos
- import_expansion_file

Build a Polynomial Chaos Expansion (PCE) by import coefficients and a multi-index from a file

Specification

Alias: none
Argument(s): STRING

Description

The coefficients can be specified in an arbitrary order. The multi-index provided is used to generate a sparse expansion that conists only of the indices corresponding to the non-zero coefficients provided in the file.

variance_based_decomp

- Keywords Area
- method
- polynomial_chaos
- variance_based_decomp

Global sensitivity analyis using variance-based decomposition

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		interaction_order	Specify the maximum number of variables allowed in an interaction when reporting interaction metrics.
	Optional		drop_tolerance	Drop values lower than this tolerance

Description

Dakota can calculate sensitivity indices through variance based decomposition using the keyword `variance_based_decomp`. These indicate how important the uncertainty in each input variable is in contributing to the output variance.

`variance_based_decomp` is extremely computationally intensive since replicated sets of sample values are evaluated. If the user specified a number of samples, N, and a number of nondeterministic variables, M, variance-based decomposition requires the evaluation of N∗(M+2) samples. **Note that specifying this keyword will increase the number of function evaluations above the number requested with the `samples` keyword** .

Because of the computational cost, `variance_based_decomp` is turned off as a default.

To obtain sensitivity indices that are reasonably accurate, we recommend that N, the number of samples, be at least one hundred and preferably several hundred or thousands.

Theory

In this context, we take sensitivity analysis to be global, not local as when calculating derivatives of output variables with respect to input variables. Our definition is similar to that of [[74] "Saltelli et al., 2004"]: "The study of how uncertainty in the output of a model can be apportioned to different sources of uncertainty in the model input."

Variance based decomposition is a way of using sets of samples to understand how the variance of the output behaves, with respect to each input variable. A larger value of the sensitivity index, S_i, means that the uncertainty in the input variable i has a larger effect on the variance of the output. More details on the calculations and interpretation of the sensitivity indices can be found in [[74] "Saltelli et al., 2004"] and [[88] "Weirs et al., 2010"].

interaction_order

- Keywords Area
- method
- polynomial_chaos
- variance_based_decomp
- interaction_order

Specify the maximum number of variables allowed in an interaction when reporting interaction metrics.

Specification

Alias: none
Argument(s): INTEGER

Description

The interaction_order option has been added to allow suppression of higher-order interactions, since the output volume (and memory and compute consumption) of these results could be extensive for high dimensional problems (note: the previous univariate_effects specification is equivalent to interaction_order = 1 in the current specification). Similar to suppression of interactions is the covariance control, which can be selected to be diagonal_covariance or full_covariance, with the former supporting suppression of the off-diagonal covariance terms (to again save compute and memory resources and reduce output volume)

drop_tolerance

- Keywords Area
- method
- polynomial_chaos
- variance_based_decomp
- drop_tolerance

Drop values lower than this tolerance

Specification

Alias: none
Argument(s): REAL

Description

Drop values lower than this tolerance. For PCE which outputs main, interaction, and total effects by default, the univariate_effects option has been added to allow suppression of the interaction effects since the output volume of these results can be prohibitive for high dimensional problems. Similar to suppression of these interactions is the covariance control, which can be selected to be diagonal_covariance or full_covariance, with the former supporting suppression of the off-diagonal covariance terms (to save compute and memory resources and reduce output volume).

diagonal_covariance

- Keywords Area
- method
- polynomial_chaos
- diagonal_covariance

Display only the diagonal terms of the covariance matrix

Specification

Alias: none
Argument(s): none

Description

With a large number of responses, the covariance matrix can be very large. `diagonal_covariance` is used to suppress the off-diagonal covariance terms (to save compute and memory resources and reduce output volume).

full_covariance

- Keywords Area
- method
- polynomial_chaos
- full_covariance

Display the full covariance matrix

Specification

Alias: none
Argument(s): none

Description

With a large number of responses, the covariance matrix can be very large. `full_covariance` is used to force Dakota to output the full covariance matrix.

normalized

- Keywords Area
- method
- polynomial_chaos
- normalized

The normalized specification requests output of PCE coefficients that correspond to normalized orthogonal basis polynomials

Specification

Alias: none
Argument(s): none

Description

The normalized specification requests output of PCE coefficients that correspond to normalized orthogonal basis polynomials

sample_type

- Keywords Area
- method
- polynomial_chaos
- sample_type

Choose the sampling algorithm

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	lhs	Use LHS method
			random	Purely random samples

Description

There are only two types of sampling:

- Monte Carlo (pure random)
- Latin Hypercube Sampling

However, some methods allow an existing study to be augmented with additional or incremental sampling. This is used to get better estimates of mean, variance, and percentiles. The incremental study must use the same type of sampling as the existing study.

lhs

- Keywords Area
- method
- polynomial_chaos
- sample_type
- lhs

Use LHS method

Specification

Alias: none
Argument(s): none

Description

Use Latin Hypercube Sampling

random

- Keywords Area
- method
- polynomial_chaos
- sample_type
- random

Purely random samples

Specification

Alias: none
Argument(s): none

Description

Purely random samples.

probability_refinement

- Keywords Area
- method
- polynomial_chaos
- probability_refinement

Allow refinement of probability and generalized reliability results using importance sampling

Topics

This keyword is related to the topics:

- problem

Specification

Alias: sample_refinement
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	import	Sampling option
			adapt_import	Importance sampling option
			mm_adapt_import	Sampling option
	Optional		refinement_-samples	Specify the number of samples used to improve a probabilty estimate.

Description

Allow refinement of probability and generalized reliability results using importance sampling

The `probability_refinement` density reweighting must account for additional non-Gaussian cases. This option allows for refinement of probability and generalized reliability results using importance sampling.

import

- Keywords Area
- method
- polynomial_chaos
- probability_refinement
- import

Sampling option

Specification

Alias: none

Argument(s): none

Description

`import` centers a sampling density at one of the initial LHS samples identified in the failure region. It then generates the importance samples, weights them by their probability of occurence given the original density, and calculates the required probability (CDF or CCDF level).

adapt_import

- Keywords Area
- method
- polynomial_chaos
- probability_refinement
- adapt_import

Importance sampling option

Specification

Alias: none
Argument(s): none

Description

`adapt_import` centers a sampling density at one of the initial LHS samples identified in the failure region. It then generates the importance samples, weights them by their probability of occurence given the original density, and calculates the required probability (CDF or CCDF level). This continues iteratively until the failure probability estimate converges.

mm_adapt_import

- Keywords Area
- method
- polynomial_chaos
- probability_refinement
- mm_adapt_import

Sampling option

Specification

Alias: none
Argument(s): none

Description

`mm_adapt_import` starts with all of the samples located in the failure region to build a multimodal sampling density. First, it uses a small number of samples around each of the initial samples in the failure region. Note that these samples are allocated to the different points based on their relative probabilities of occurrence: more probable points get more samples. This early part of the approach is done to search for ”representative” points. Once these are located, the multimodal sampling density is set and then `mm_adapt_import` proceeds similarly to `adapt_import` (sample until convergence).

refinement_samples

- Keywords Area
- method
- polynomial_chaos
- probability_refinement
- refinement_samples

Specify the number of samples used to improve a probabilty estimate.

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples used to improve a probabilty estimate. If using uni-modal sampling all samples are assigned to the sampling center. If using multi-modal sampling the samples are split between mutiple samples according to some internally computed weights.

export_points_file

- Keywords Area
- method
- polynomial_chaos
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- polynomial_chaos
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- polynomial_chaos
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_expansion_file

- Keywords Area
- method
- polynomial_chaos
- export_expansion_file

Export the coefficients and multi-index of a Polynomial Chaos Expansion (PCE) to a file

Specification

Alias: none
Argument(s): STRING

Description

Export the coefficients and multi-index of a Polynomial Chaos Expansion (PCE) to a file. The multi-index written will be sparse. Spcifically the expansion will conists only of the indices corresponding to the non-zero coefficients of the PCE.

fixed_seed

- Keywords Area
- method
- polynomial_chaos
- fixed_seed

Reuse the same seed value for multiple sampling sets

Specification

Alias: none
Argument(s): none

Description

The `fixed_seed` flag is relevant if multiple sampling sets will be generated. This occurs when using advance methods (e.g., surrogate-based optimization, optimization under uncertainty).

Specifying this keyword results in the reuse of the same seed value for each of these multiple sampling sets, which can be important for reducing variability in the sampling results.

However, this behavior is not the default as the repetition of the same sampling pattern can result in a modeling weakness that an optimizer could potentially exploit (resulting in actual reliabilities that are lower than the estimated reliabilities). In either case (`fixed_seed` or not), the study is repeatable if the user specifies a `seed` and the study is random is the user omits a `seed` specification.

reliability_levels

- Keywords Area
- method
- polynomial_chaos
- reliability_levels

Specify reliability levels at which the response values will be estimated

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_reliability_-levels	Specify which `reliability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF reliabilities by projecting out the prescribed number of sample standard deviations from the sample mean.

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_reliability_levels

- Keywords Area
- method
- polynomial_chaos
- reliability_levels
- num_reliability_levels

Specify which `reliability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

response_levels

- Keywords Area
- method
- polynomial_chaos
- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_-levels` correspond to which response

	Optional	compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- polynomial_chaos
- response_levels

- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- polynomial_chaos
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none
Argument(s): none

	Required/-Optional	**Description of Group**	**Dakota Keyword**	**Dakota Keyword Description**
	Required*(Choose One)*	**Group 1**	probabilities	Compute probabilities
			reliabilities	Compute reliabilities
			gen_reliabilities	Compute generalized reliabilities
	Optional		system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- polynomial_chaos
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

reliabilities

- Keywords Area
- method
- polynomial_chaos
- response_levels
- compute
- reliabilities

Compute reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- polynomial_chaos
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- polynomial_chaos
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	series	?????
			parallel	?????

series

- Keywords Area
- method
- polynomial_chaos
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

parallel

- Keywords Area
- method
- polynomial_chaos
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

distribution

- Keywords Area
- method
- polynomial_chaos
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	cumulative	Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.

This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- polynomial_chaos
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- polynomial_chaos
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- polynomial_chaos
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- polynomial_chaos
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- polynomial_chaos
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- polynomial_chaos
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

rng

- Keywords Area
- method
- polynomial_chaos
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** mt19937	**Dakota Keyword Description** Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- polynomial_chaos
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: http://en.wikipedia.org/wiki/Mersenne_twister

rnum2

- Keywords Area
- method
- polynomial_chaos
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

samples

- Keywords Area
- method
- polynomial_chaos
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- polynomial_chaos
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- polynomial_chaos
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.52 stoch_collocation

- Keywords Area
- method
- stoch_collocation

Uncertainty quantification with stochastic collocation

Specification

Alias: nond_stoch_collocation
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional(*Choose One*)	**Group 1**	p_refinement	Automatic polynomial order refinement
			h_refinement	Employ h-refinement to refine the grid
	Optional(*Choose One*)	**Group 2**	piecewise	Use piecewise local basis functions
			askey	Select the Askey basis polynomials that match the basis random variable

			wiener	Use Hermite polynomial basis functions
	Required*(Choose One)*	**Group 3**	quadrature_order	Cubature using tensor-products of Gaussian quadrature rules
			sparse_grid_level	Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids
	Optional		dimension_-preference	A set of weights specifying the realtive importance of each uncertain variable (dimension)
	Optional		use_derivatives	Use derivative data to construct surrogate models
	Optional*(Choose One)*	**Group 4**	nested	Enforce use of nested quadrature rules if available
			non_nested	Enforce use of non-nested quadrature rules
	Optional		variance_based_-decomp	Global sensitivity analyis using variance-based decomposition
	Optional*(Choose One)*	**Group 5**	diagonal_-covariance	Display only the diagonal terms of the covariance matrix
			full_covariance	Display the full covariance matrix
	Optional		sample_type	Choose the sampling algorithm
	Optional		probability_-refinement	Allow refinement of probability and generalized reliability results using importance sampling

	Optional	export_points_file	Specify file to write out the responses from the surrogate
	Optional	fixed_seed	Reuse the same seed value for multiple sampling sets
	Optional	reliability_levels	Specify reliability levels at which the response values will be estimated
	Optional	response_levels	Specify the response values at which to estimate relevant statistics
	Optional	distribution	Select between cumulative or complementary-cumulative functions
	Optional	probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional	gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional	rng	Choose a random number generator
	Optional	samples	Specify the number of samples
	Optional	seed	Specify the seed of the random number generator
	Optional	model_pointer	Specify which model block will be used by a method

Description

Stochastic collocation is a general framework for approximate representation of random response functions in terms of finite-dimensional interpolation bases. The following provides details on the various stochastic collocation method options in Dakota.

The groups and optional keywords relating to method specification are:

- Group 1
- Group 2
- Group 3
- Group 4
- dimension_preference
- use_derivatives
- fixed_seed

In addition, this method treats variables that are not aleatoric-uncertain different, despite the active keyword.

Group 5 and the remainder of the optional keywords relate to the output of the method.

stoch_collocation Group 2

SC supports four types of bases:

the option of piecewise local basis functions. These are piecewise linear splines, or in the case of gradient-enhanced interpolation via the use_derivatives specification, piecewise cubic Hermite splines. Both of these basis selections provide local support only over the range from the interpolated point to its nearest 1D neighbors (within a tensor grid or within each of the tensor grids underlying a sparse grid), which exchanges the fast convergence of global bases for smooth functions for robustness in the representation of nonsmooth response functions (that can induce Gibbs oscillations when using high-order global basis functions). When local basis functions are used, the usage of nonequidistant collocation points (e.g., the Gauss point selections described above) is not well motivated, so equidistant Newton-Cotes points are employed in this case, and all random variable types are transformed to standard uniform probability space. The global gradient-enhanced interpolants (Hermite interpolation polynomials) are also restricted to uniform or transformed uniform random variables (due to the need to compute collocation weights by integration of the basis polynomials) and share the variable support shown in these tables for Piecewise SE. Due to numerical instability in these high-order basis polynomials, they are deactivated by default but can be activated by developers using a compile-time switch.

Another distinguishing characteristic of stochastic collocation relative to polynomial chaos is the ability to reformulate the interpolation problem from a nodal interpolation approach into a hierarchical formulation in which each new level of interpolation defines a set of incremental refinements (known as hierarchical surpluses) layered on top of the interpolants from previous levels. This formulation lends itself naturally to uniform or adaptive refinement strategies, since the hierarchical surpluses can be interpreted as error estimates for the interpolant. Either global or local/piecewise interpolants in either value-based or gradient-enhanced approaches can be formulated using hierarchical interpolation. The primary restriction for the hierarchical case is that it currently requires a sparse grid approach using nested quadrature rules (Genz-Keister, Gauss-Patterson, or Newton-Cotes for standard normals and standard uniforms in a transformed space: Askey, Wiener, or Piecewise settings may be required), although this restriction can be relaxed in the future. A selection of hierarchical interpolation will provide greater precision in the increments to mean, standard deviation, covariance, and reliability-based level mappings induced by a grid change within uniform or goal-oriented adaptive refinement approaches (see following section).

Automated expansion refinement can be selected as either p_refinement or h_refinement, and either refinement specification can be either uniform or dimension_adaptive. The dimension_adaptive case can be further specified as either sobol or generalized (decay not supported). Each of these automated refinement approaches makes use of the max_iterations and convergence_tolerance iteration controls. The h_refinement specification involves use of the same piecewise interpolants (linear or cubic Hermite splines) described above for the piecewise specification option (it is not necessary to redundantly specify piecewise in the case of h_refinement). In future releases, the hierarchical interpolation approach will enable local refinement in addition to the current uniform and dimension_adaptive options.

The variance_based_decomp and covariance controls. Interpolation points for these dimensions are based on Gauss-Legendre rules if non-nested, Gauss-Patterson rules if nested, and Newton-Cotes points in the case of piecewise bases. Again, when probability integrals are evaluated, only the aleatory random variable domain is integrated, leaving behind a polynomial relationship between the statistics and the remaining design/state/epistemic variables.

To form the multidimensional interpolants L_i of the expansion, two options are provided.

1. interpolation on a tensor-product of Gaussian quadrature points (specified with quadrature_order and, optionally, dimension_preference for anisotropic tensor grids). As for PCE, non-nested Gauss rules are employed by default, although the presence of p_refinement or h_refinement will result in default usage of nested rules for normal or uniform variables after any variable transformations have been applied (both defaults can be overridden using explicit nested or non_nested specifications).

2. interpolation on a Smolyak sparse grid (specified with sparse_grid_level and, optionally, dimension_preference for anisotropic sparse grids) defined from Gaussian rules. As for sparse PCE, nested rules are employed unless overridden with the non_nested option, and the growth rules are restricted unless overridden by the unrestricted keyword.

If n is small, then tensor-product Gaussian quadrature is again the preferred choice. For larger n, tensor-product quadrature quickly becomes too expensive and the sparse grid approach is preferred. For self-consistency in growth rates, nested rules employ restricted exponential growth (with the exception of the dimension_adaptive p_refinement generalized case) for consistency with the linear growth used for non-nested Gauss rules (integrand precision $i = 4l + 1$ for sparse grid level l and $i = 2m - 1$ for tensor grid order m).

stoch_collocation Group 5

These two keywords are used to specify how this method outputs the covariance of the responses.

Optional Keywords regarding method outputs

Each of these sampling specifications refer to sampling on the PCE approximation for the purposes of generating approximate statistics.

- sample_type
- samples
- seed
- fixed_seed
- rng
- probability_refinement
- reliability_levels
- response_levels
- probability_levels
- gen_reliability_levels

which should be distinguished from simulation sampling for generating the PCE coefficients as described in options 4 and 5 above (although options 4 and 5 will share the sample_type, seed, and rng settings, if provided).

When using the probability_refinement control, the number of refinement samples is not under the user's control (these evaluations are approximation-based, so management of this expense is less critical). This option allows for refinement of probability and generalized reliability results using importance sampling.

Theory

The stochastic collocation (SC) method is very similar to the PCE method, with the key difference that the orthogonal polynomial basis functions are replaced with interpolation polynomial bases. The interpolation polynomials may be either local or global and either value-based or gradient-enhanced. In the local case, valued-based are piecewise linear splines and gradient-enhanced are piecewise cubic splines, and in the global case, valued-based are Lagrange interpolants and gradient-enhanced are Hermite interpolants. A value-based expansion takes the form

$$R = \sum_{i=1}^{N_p} r_i L_i(\xi)$$

where N_p is the total number of collocation points, r_i is a response value at the ith collocation point, L_i is the ith multidimensional interpolation polynomial, and ξ is a vector of standardized random variables. The ith interpolation polynomial assumes the value of 1 at the ith collocation point and 0 at all other collocation points, involving either a global Lagrange polynomial basis or local piecewise splines. It is easy to see that the approximation reproduces the response values at the collocation points and interpolates between these values at other points. A gradient-enhanced expansion (selected via the `use_derivatives` keyword) involves both type 1 and type 2 basis functions as follows:

$$R = \sum_{i=1}^{N_p} [r_i H_i^{(1)}(\xi) + \sum_{j=1}^{n} \frac{dr_i}{d\xi_j} H_{ij}^{(2)}(\xi)]$$

where the ith type 1 interpolant produces 1 for the value at the ith collocation point, 0 for values at all other collocation points, and 0 for derivatives (when differentiated) at all collocation points, and the ij^{th} type 2 interpolant produces 0 for values at all collocation points, 1 for the jth derivative component at the ith collocation point, and 0 for the jth derivative component at all other collocation points. Again, this expansion reproduces the response values at each of the collocation points, and when differentiated, also reproduces each component of the gradient at each of the collocation points. Since this technique includes the derivative interpolation explicitly, it eliminates issues with matrix ill-conditioning that can occur in the gradient-enhanced PCE approach based on regression. However, the calculation of high-order global polynomials with the desired interpolation properties can be similarly numerically challenging such that the use of local cubic splines is recommended due to numerical stability.

Thus, in PCE, one forms coefficients for known orthogonal polynomial basis functions, whereas SC forms multidimensional interpolation functions for known coefficients. Dakota provides access to SC methods through the NonDStochCollocation class. Refer to the Uncertainty Quantification Capabilities chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information on the SC algorithm.

When using multifidelity UQ, the high fidelity expansion generated from combining the low fidelity and discrepancy expansions retains the polynomial form of the low fidelity expansion (only the coefficients are updated).

See Also

These keywords may also be of interest:

- adaptive_sampling
- gpais
- local_reliability
- global_reliability

- sampling
- importance_sampling
- polynomial_chaos

p_refinement

- Keywords Area
- method
- stoch_collocation
- p_refinement

Automatic polynomial order refinement

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** uniform	**Dakota Keyword Description** Refine an expansion uniformly in all dimensions.
			dimension_-adaptive	Perform automated expansion refinement by adapting to dimensions with higher 'importance'.

Description

The `p_refinement` keyword specifies the usage of automated polynomial order refinement, which can be either `uniform` or `dimension_adaptive`.

The `dimension_adaptive` option is supported for the tensor-product quadrature and Smolyak sparse grid options and `uniform` is supported for tensor and sparse grids as well as regression approaches (`collocation_points` or `collocation_ratio`).

Each of these refinement cases makes use of the `max_iterations` and `convergence_tolerance` method independent controls. The former control limits the number of refinement iterations, and the latter control terminates refinement when the two-norm of the change in the response covariance matrix (or, in goal-oriented approaches, the two-norm of change in the statistical quantities of interest (QOI)) falls below the tolerance.

The `dimension_adaptive` case can be further specified to utilize `sobol`, `decay`, or `generalized` refinement controls. The former two cases employ anisotropic tensor/sparse grids in which the anisotropic dimension preference (leading to anisotropic integrations/expansions with differing refinement levels for different random dimensions) is determined using either total Sobol' indices from variance-based decomposition (`sobol` case: high indices result in high dimension preference) or using spectral coefficient decay rates from a rate

estimation technique similar to Richardson extrapolation (`decay` case: low decay rates result in high dimension preference). In these two cases as well as the `uniform` refinement case, the `quadrature_order` or `sparse_grid_level` are ramped by one on each refinement iteration until either of the two convergence controls is satisfied. For the `uniform` refinement case with regression approaches, the `expansion_order` is ramped by one on each iteration while the oversampling ratio (either defined by `collocation_ratio` or inferred from `collocation_points` based on the initial expansion) is held fixed. Finally, the `generalized` `dimension_adaptive` case is the default adaptive approach; it refers to the generalized sparse grid algorithm, a greedy approach in which candidate index sets are evaluated for their impact on the statistical QOI, the most influential sets are selected and used to generate additional candidates, and the index set frontier of a sparse grid is evolved in an unstructured and goal-oriented manner (refer to User's Manual PCE descriptions for additional specifics).

For the case of p_refinement or the case of an explicit nested override, Gauss-Hermite rules are replaced with Genz-Keister nested rules and Gauss-Legendre rules are replaced with Gauss-Patterson nested rules, both of which exchange lower integrand precision for greater point reuse.

uniform

- Keywords Area
- method
- stoch_collocation
- p_refinement
- uniform

Refine an expansion uniformly in all dimensions.

Specification

Alias: none
Argument(s): none

Description

The quadrature_order or sparse_grid_level are ramped by one on each refinement iteration until either of the two convergence controls is satisfied. For the uniform refinement case with regression approaches, the expansion_order is ramped by one on each iteration while the oversampling ratio (either defined by collocation_ratio or inferred from collocation_points based on the initial expansion) is held fixed.

dimension_adaptive

- Keywords Area
- method
- stoch_collocation
- p_refinement
- dimension_adaptive

Perform automated expansion refinement by adapting to dimensions with higher 'importance'.

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	sobol	Determine dimension preference for refinement of a polynomial chaos expansion from the sobol sensitivity indices of the PCE.
			generalized	Use the generalized sparse grid dimension adaptive algorithm to refine a sparse grid approximation of polynomial expansion.

Description

Perform automated expansion refinement by adapting to dimensions with higher 'importance'. Dimension_-adaptive refinement can be further specified as either sobol or generalized (decay not supported). Each of these automated refinement approaches makes use of the max_iterations and convergence_tolerance iteration controls

sobol

- Keywords Area
- method
- stoch_collocation
- p_refinement
- dimension_adaptive
- sobol

Determine dimension preference for refinement of a polynomial chaos expansion from the sobol sensitivity indices of the PCE.

Specification

Alias: none

Argument(s): none

Description

Determine dimension preference for refinement of a polynomial chaos expansion from the sobol sensitivity indices of the PCE. High indices result in high dimension preference.

generalized

- Keywords Area
- method
- stoch_collocation
- p_refinement
- dimension_adaptive
- generalized

Use the generalized sparse grid dimension adaptive algorithm to refine a sparse grid approximation of polynomial expansion.

Specification

Alias: none
Argument(s): none

Description

The generalized sparse grid algorithm is a greedy approach in which candidate index sets are evaluated for their impact on the statistical QOI, the most influential sets are selected and used to generate additional candidates, and the index set frontier of a sparse grid is evolved in an unstructured and goal-oriented manner (refer to User's Manual PCE descriptions for additional specifics).

h_refinement

- Keywords Area
- method
- stoch_collocation
- h_refinement

Employ h-refinement to refine the grid

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	Description of Group **Group 1**	Dakota Keyword uniform	Dakota Keyword Description Refine an expansion uniformly in all dimensions.
			dimension_-adaptive	Select the adaptive Metropolis sampler
			local_adaptive	Select the adaptive Metropolis sampler

Description

Automated expansion refinement can be selected as either p_refinement or h_refinement, and either refinement specification can be either uniform or dimension_adaptive. The dimension_adaptive case can be further specified as either sobol or generalized (decay not supported). Each of these automated refinement approaches makes use of the max_iterations and convergence_tolerance iteration controls. The h_refinement specification involves use of the same piecewise interpolants (linear or cubic Hermite splines) described above for the piecewise specification option (it is not necessary to redundantly specify piecewise in the case of h_refinement). In future releases, the hierarchical interpolation approach will enable local refinement in addition to the current uniform and dimension-_adaptive options.

uniform

- Keywords Area
- method
- stoch_collocation
- h_refinement
- uniform

Refine an expansion uniformly in all dimensions.

Specification

Alias: none
Argument(s): none

Description

The quadrature_order or sparse_grid_level are ramped by one on each refinement iteration until either of the two convergence controls is satisfied. For the uniform refinement case with regression approaches, the expansion-_order is ramped by one on each iteration while the oversampling ratio (either defined by collocation_ratio or inferred from collocation_points based on the initial expansion) is held fixed.

dimension_adaptive

- Keywords Area
- method
- stoch_collocation
- h_refinement
- dimension_adaptive

Select the adaptive Metropolis sampler

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required(*Choose One*)	**Group 1**	sobol	Determine dimension preference for refinement of a polynomial chaos expansion from the sobol sensitivity indices of the PCE.
			generalized	Use the generalized sparse grid dimension adaptive algorithm to refine a sparse grid approximation of polynomial expansion.

Description

The use of an adaptive Metropolis algorithm in the Markov Chain Monte Carlo algorithm, involving adaptive updates to the covariance matrix.

sobol

- Keywords Area
- method
- stoch_collocation
- h_refinement
- dimension_adaptive

- sobol

Determine dimension preference for refinement of a polynomial chaos expansion from the sobol sensitivity indices of the PCE.

Specification

Alias: none
Argument(s): none

Description

Determine dimension preference for refinement of a polynomial chaos expansion from the sobol sensitivity indices of the PCE. High indices result in high dimension preference.

generalized

- Keywords Area
- method
- stoch_collocation
- h_refinement
- dimension_adaptive
- generalized

Use the generalized sparse grid dimension adaptive algorithm to refine a sparse grid approximation of polynomial expansion.

Specification

Alias: none
Argument(s): none

Description

The generalized sparse grid algorithm is a greedy approach in which candidate index sets are evaluated for their impact on the statistical QOI, the most influential sets are selected and used to generate additional candidates, and the index set frontier of a sparse grid is evolved in an unstructured and goal-oriented manner (refer to User's Manual PCE descriptions for additional specifics).

local_adaptive

- Keywords Area
- method
- stoch_collocation
- h_refinement
- local_adaptive

Select the adaptive Metropolis sampler

Specification

Alias: none
Argument(s): none

Description

The use of an adaptive Metropolis algorithm in the Markov Chain Monte Carlo algorithm, involving adaptive updates to the covariance matrix.

piecewise

- Keywords Area
- method
- stoch_collocation
- piecewise

Use piecewise local basis functions

Specification

Alias: none
Argument(s): none

Description

SC also supports the option of piecewise local basis functions. These are piecewise linear splines, or in the case of gradient-enhanced interpolation via the use_derivatives specification, piecewise cubic Hermite splines. Both of these basis selections provide local support only over the range from the interpolated point to its nearest 1D neighbors (within a tensor grid or within each of the tensor grids underlying a sparse grid), which exchanges the fast convergence of global bases for smooth functions for robustness in the representation of nonsmooth response functions (that can induce Gibbs oscillations when using high-order global basis functions). When local basis functions are used, the usage of nonequidistant collocation points (e.g., the Gauss point selections described above) is not well motivated, so equidistant Newton-Cotes points are employed in this case, and all random variable types are transformed to standard uniform probability space.

askey

- Keywords Area
- method
- stoch_collocation
- askey

Select the Askey basis polynomials that match the basis random variable

Specification

Alias: none
Argument(s): none

Description

The Askey option employs an extended basis of Hermite, Legendre, Laguerre, Jacobi, and generalized Laguerre orthogonal polynomials.

- Askey polynomial selections include Hermite for normal (optimal) as well as bounded normal, lognormal, bounded lognormal, gumbel, frechet, and weibull (sub-optimal); Legendre for uniform (optimal) as well as loguniform, triangular, and bin-based histogram (sub-optimal); Laguerre for exponential (optimal); Jacobi for beta (optimal); and generalized Laguerre for gamma (optimal).

See Also

These keywords may also be of interest:

- wiener

wiener

- Keywords Area
- method
- stoch_collocation
- wiener

Use Hermite polynomial basis functions

Specification

Alias: none
Argument(s): none

Description

The Wiener option uses a Hermite orthogonal polynomial basis for all random variables and employs the same nonlinear variable transformation as the local and global reliability methods (and therefore has the same variable support).

See Also

These keywords may also be of interest:

- askey

quadrature_order

- Keywords Area
- method
- stoch_collocation
- quadrature_order

Cubature using tensor-products of Gaussian quadrature rules

Specification

Alias: none
Argument(s): INTEGERLIST

Description

Multidimensional integration by a tensor-product of Gaussian quadrature rules (specified with `quadrature_order`, and, optionally, `dimension_preference`). The default rule selection is to employ `non_nested` Gauss rules including Gauss-Hermite (for normals or transformed normals), Gauss-Legendre (for uniforms or transformed uniforms), Gauss-Jacobi (for betas), Gauss-Laguerre (for exponentials), generalized Gauss-Laguerre (for gammas), and numerically-generated Gauss rules (for other distributions when using an Extended basis). For the case of `p_refinement` or the case of an explicit `nested` override, Gauss-Hermite rules are replaced with Genz-Keister nested rules and Gauss-Legendre rules are replaced with Gauss-Patterson nested rules, both of which exchange lower integrand precision for greater point reuse. By specifying a `dimension_preference`, where higher preference leads to higher order polynomial resolution, the tensor grid may be rendered anisotropic. The dimension specified to have highest preference will be set to the specified `quadrature_order` and all other dimensions will be reduced in proportion to their reduced preference; any non-integral portion is truncated. To synchronize with tensor-product integration, a tensor-product expansion is used, where the order p_i of the expansion in each dimension is selected to be half of the integrand precision available from the rule in use, rounded down. In the case of non-nested Gauss rules with integrand precision $2m_i - 1$, p_i is one less than the quadrature order m_i in each dimension (a one-dimensional expansion contains the same number of terms, $p + 1$, as the number of Gauss points). The total number of terms, *N*, in a tensor-product expansion involving *n* uncertain input variables is

$$N = 1 + P = \prod_{i=1}^{n} (p_i + 1)$$

In some advanced use cases (e.g., multifidelity UQ), multiple grid resolutions can be employed; for this reason, the `quadrature_order` specification supports an array input.

sparse_grid_level

- Keywords Area
- method
- stoch_collocation
- sparse_grid_level

Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Specification

Alias: none

Argument(s): INTEGERLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	restricted	Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids
			unrestricted	Overide the default restriction of growth rates for nested and non-nested rules that are by defualt synchronized for consistency.
	Optional*(Choose One)*	**Group 2**	nodal	Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids
			hierarchical	Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Description

Multi-dimensional integration by the Smolyak sparse grid method (specified with sparse_grid_level and, optionally, dimension_preference). The underlying one-dimensional integration rules are the same as for the tensor-product quadrature case; however, the default rule selection is nested for sparse grids (Genz-Keister for normals/transformed normals and Gauss-Patterson for uniforms/transformed uniforms). This default can be overridden with an explicit non_nested specification (resulting in Gauss-Hermite for normals/transformed normals and Gauss-Legendre for uniforms/transformed uniforms). As for tensor quadrature, the dimension_preference specification enables the use of anisotropic sparse grids (refer to the PCE description in the User's Manual for the anisotropic index set constraint definition). Similar to anisotropic tensor grids, the dimension with greatest preference will have resolution at the full sparse_grid_level and all other dimension resolutions will be reduced in proportion to their reduced preference. For PCE with either isotropic or anisotropic sparse grids, a summation of tensor-product expansions is used, where each anisotropic tensor-product quadrature rule underlying the sparse grid construction results in its own anisotropic tensor-product expansion as described in case 1. These anisotropic

tensor-product expansions are summed into a sparse PCE using the standard Smolyak summation (again, refer to the User's Manual for additional details). As for quadrature_order, the sparse_grid_level specification admits an array input for enabling specification of multiple grid resolutions used by certain advanced solution methodologies.

This keyword can be used when using sparse grid integration to calculate PCE coefficients or when generating samples for sparse grid collocation.

restricted

- Keywords Area
- method
- stoch_collocation
- sparse_grid_level
- restricted

Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Specification

Alias: none

Argument(s): none

Description

Multi-dimensional integration by the Smolyak sparse grid method (specified with sparse_grid_level and, optionally, dimension_preference). The underlying one-dimensional integration rules are the same as for the tensor-product quadrature case; however, the default rule selection is nested for sparse grids (Genz-Keister for normals/transformed normals and Gauss-Patterson for uniforms/transformed uniforms). This default can be overridden with an explicit non_nested specification (resulting in Gauss-Hermite for normals/transformed normals and Gauss-Legendre for uniforms/transformed uniforms). As for tensor quadrature, the dimension_preference specification enables the use of anisotropic sparse grids (refer to the PCE description in the User's Manual for the anisotropic index set constraint definition). Similar to anisotropic tensor grids, the dimension with greatest preference will have resolution at the full sparse_grid_level and all other dimension resolutions will be reduced in proportion to their reduced preference. For PCE with either isotropic or anisotropic sparse grids, a summation of tensor-product expansions is used, where each anisotropic tensor-product quadrature rule underlying the sparse grid construction results in its own anisotropic tensor-product expansion as described in case 1. These anisotropic tensor-product expansions are summed into a sparse PCE using the standard Smolyak summation (again, refer to the User's Manual for additional details). As for quadrature_order, the sparse_grid_level specification admits an array input for enabling specification of multiple grid resolutions used by certain advanced solution methodologies.

This keyword can be used when using sparse grid integration to calculate PCE coefficients or when generating samples for sparse grid collocation.

unrestricted

- Keywords Area
- method
- stoch_collocation
- sparse_grid_level
- unrestricted

Overide the default restriction of growth rates for nested and non-nested rules that are by defualt synchronized for consistency.

Specification

Alias: none
Argument(s): none

Description

In the quadrature and sparse grid cases, growth rates for nested and non-nested rules can be synchronized for consistency. For a non-nested Gauss rule used within a sparse grid, linear one-dimensional growth rules of $m = 2l + 1$ are used to enforce odd quadrature orders, where l is the grid level and m is the number of points in the rule. The precision of this Gauss rule is then $i = 2m - 1 = 4l + 1$. For nested rules, order growth with level is typically exponential; however, the default behavior is to restrict the number of points to be the lowest order rule that is available that meets the one-dimensional precision requirement implied by either a level l for a sparse grid ($i = 4l + 1$) or an order m for a tensor grid ($i = 2m - 1$). This behavior is known as "restricted growth" or "delayed sequences." To override this default behavior in the case of sparse grids, the unrestricted keyword can be used; it cannot be overridden for tensor grids using nested rules since it also provides a mapping to the available nested rule quadrature orders. An exception to the default usage of restricted growth is the dimension_adaptive p-_refinement generalized sparse grid case described previously, since the ability to evolve the index sets of a sparse grid in an unstructured manner eliminates the motivation for restricting the exponential growth of nested rules.

nodal

- Keywords Area
- method
- stoch_collocation
- sparse_grid_level
- nodal

Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Specification

Alias: none
Argument(s): none

Description

Multi-dimensional integration by the Smolyak sparse grid method (specified with sparse_grid_level and, optionally, dimension_preference). The underlying one-dimensional integration rules are the same as for the tensor-product quadrature case; however, the default rule selection is nested for sparse grids (Genz-Keister for normals/transformed normals and Gauss-Patterson for uniforms/transformed uniforms). This default can be overridden with an explicit non_nested specification (resulting in Gauss-Hermite for normals/transformed normals and Gauss-Legendre for uniforms/transformed uniforms). As for tensor quadrature, the dimension_preference specification enables the use of anisotropic sparse grids (refer to the PCE description in the User's Manual for the anisotropic index set constraint definition). Similar to anisotropic tensor grids, the dimension with greatest preference will have resolution at the full sparse_grid_level and all other dimension resolutions will be reduced in proportion to their reduced preference. For PCE with either isotropic or anisotropic sparse grids, a summation of tensor-product expansions is used, where each anisotropic tensor-product quadrature rule underlying the sparse grid construction results in its own anisotropic tensor-product expansion as described in case 1. These anisotropic tensor-product expansions are summed into a sparse PCE using the standard Smolyak summation (again, refer to the User's Manual for additional details). As for quadrature_order, the sparse_grid_level specification admits an array input for enabling specification of multiple grid resolutions used by certain advanced solution methodologies.

This keyword can be used when using sparse grid integration to calculate PCE coefficients or when generating samples for sparse grid collocation.

hierarchical

- Keywords Area
- method
- stoch_collocation
- sparse_grid_level
- hierarchical

Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Specification

Alias: none

Argument(s): none

Description

Multi-dimensional integration by the Smolyak sparse grid method (specified with sparse_grid_level and, optionally, dimension_preference). The underlying one-dimensional integration rules are the same as for the tensor-product quadrature case; however, the default rule selection is nested for sparse grids (Genz-Keister for normals/transformed normals and Gauss-Patterson for uniforms/transformed uniforms). This default can be overridden with an explicit non_nested specification (resulting in Gauss-Hermite for normals/transformed normals and Gauss-Legendre for uniforms/transformed uniforms). As for tensor quadrature, the dimension_preference specification enables the use of anisotropic sparse grids (refer to the PCE description in the User's Manual for the anisotropic index set constraint definition). Similar to anisotropic tensor grids, the dimension with greatest preference will have resolution at the full sparse_grid_level and all other dimension resolutions will be reduced in

proportion to their reduced preference. For PCE with either isotropic or anisotropic sparse grids, a summation of tensor-product expansions is used, where each anisotropic tensor-product quadrature rule underlying the sparse grid construction results in its own anisotropic tensor-product expansion as described in case 1. These anisotropic tensor-product expansions are summed into a sparse PCE using the standard Smolyak summation (again, refer to the User's Manual for additional details). As for quadrature_order, the sparse_grid_level specification admits an array input for enabling specification of multiple grid resolutions used by certain advanced solution methodologies.

This keyword can be used when using sparse grid integration to calculate PCE coefficients or when generating samples for sparse grid collocation.

dimension_preference

- Keywords Area
- method
- stoch_collocation
- dimension_preference

A set of weights specifying the realtive importance of each uncertain variable (dimension)

Specification

Alias: none
Argument(s): REALLIST

Description

A set of weights specifying the realtive importance of each uncertain variable (dimension). Using this specification leada to anisotropic integrations with differing refinement levels for different random dimensions.

See Also

These keywords may also be of interest:

- sobol
- decay

use_derivatives

- Keywords Area
- method
- stoch_collocation
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none
Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

nested

- Keywords Area
- method
- stoch_collocation
- nested

Enforce use of nested quadrature rules if available

Specification

Alias: none
Argument(s): none

Description

Enforce use of nested quadrature rules if available. For instance if the aleatory variables are Gaussian use the Nested Genz-Keister rule instead of the default non-nested Gauss-Hermite rule variables are

non_nested

- Keywords Area
- method
- stoch_collocation
- non_nested

Enforce use of non-nested quadrature rules

Specification

Alias: none
Argument(s): none

Description

Enforce use of non-nested quadrature rules if available. For instance if the aleatory variables are Gaussian use the non-nested Gauss-Hermite rule

variance_based_decomp

- Keywords Area
- method
- stoch_collocation
- variance_based_decomp

Global sensitivity analyis using variance-based decomposition

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		interaction_order	Specify the maximum number of variables allowed in an interaction when reporting interaction metrics.
	Optional		drop_tolerance	Drop values lower than this tolerance

Description

Dakota can calculate sensitivity indices through variance based decomposition using the keyword `variance_based_decomp`. These indicate how important the uncertainty in each input variable is in contributing to the output variance.

`variance_based_decomp` is extremely computationally intensive since replicated sets of sample values are evaluated. If the user specified a number of samples, N, and a number of nondeterministic variables, M, variance-based decomposition requires the evaluation of N∗(M+2) samples. **Note that specifying this keyword will increase the number of function evaluations above the number requested with the `samples` keyword** .

Because of the computational cost, `variance_based_decomp` is turned off as a default.

To obtain sensitivity indices that are reasonably accurate, we recommend that N, the number of samples, be at least one hundred and preferably several hundred or thousands.

Theory

In this context, we take sensitivity analysis to be global, not local as when calculating derivatives of output variables with respect to input variables. Our definition is similar to that of [[74] "Saltelli et al., 2004"]: "The study of how uncertainty in the output of a model can be apportioned to different sources of uncertainty in the model input."

Variance based decomposition is a way of using sets of samples to understand how the variance of the output behaves, with respect to each input variable. A larger value of the sensitivity index, S_i, means that the uncertainty in the input variable i has a larger effect on the variance of the output. More details on the calculations and interpretation of the sensitivity indices can be found in [[74] "Saltelli et al., 2004"] and [[88] "Weirs et al., 2010"].

interaction_order

- Keywords Area
- method
- stoch_collocation
- variance_based_decomp
- interaction_order

Specify the maximum number of variables allowed in an interaction when reporting interaction metrics.

Specification

Alias: none
Argument(s): INTEGER

Description

The interaction_order option has been added to allow suppression of higher-order interactions, since the output volume (and memory and compute consumption) of these results could be extensive for high dimensional problems (note: the previous univariate_effects specification is equivalent to interaction_order = 1 in the current specification). Similar to suppression of interactions is the covariance control, which can be selected to be diagonal_covariance or full_covariance, with the former supporting suppression of the off-diagonal covariance terms (to again save compute and memory resources and reduce output volume)

drop_tolerance

- Keywords Area
- method
- stoch_collocation
- variance_based_decomp
- drop_tolerance

Drop values lower than this tolerance

Specification

Alias: none
Argument(s): REAL

Description

Drop values lower than this tolerance. For PCE which outputs main, interaction, and total effects by default, the univariate_effects option has been added to allow suppression of the interaction effects since the output volume of these results can be prohibitive for high dimensional problems. Similar to suppression of these interactions is the covariance control, which can be selected to be diagonal_covariance or full_covariance, with the former supporting suppression of the off-diagonal covariance terms (to save compute and memory resources and reduce output volume).

diagonal_covariance

- Keywords Area
- method
- stoch_collocation
- diagonal_covariance

Display only the diagonal terms of the covariance matrix

Specification

Alias: none
Argument(s): none

Description

With a large number of responses, the covariance matrix can be very large. `diagonal_covariance` is used to suppress the off-diagonal covariance terms (to save compute and memory resources and reduce output volume).

full_covariance

- Keywords Area
- method
- stoch_collocation
- full_covariance

Display the full covariance matrix

Specification

Alias: none
Argument(s): none

Description

With a large number of responses, the covariance matrix can be very large. `full_covariance` is used to force Dakota to output the full covariance matrix.

sample_type

- Keywords Area
- method
- stoch_collocation
- sample_type

Choose the sampling algorithm

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	lhs	Use LHS method
			random	Purely random samples

Description

There are only two types of sampling:

- Monte Carlo (pure random)
- Latin Hypercube Sampling

However, some methods allow an existing study to be augmented with additional or incremental sampling. This is used to get better estimates of mean, variance, and percentiles. The incremental study must use the same type of sampling as the existing study.

lhs

- Keywords Area
- method
- stoch_collocation
- sample_type
- lhs

Use LHS method

Specification

Alias: none
Argument(s): none

Description

Use Latin Hypercube Sampling

random

- Keywords Area
- method
- stoch_collocation
- sample_type
- random

Purely random samples

Specification

Alias: none
Argument(s): none

Description

Purely random samples.

probability_refinement

- Keywords Area
- method
- stoch_collocation
- probability_refinement

Allow refinement of probability and generalized reliability results using importance sampling

Topics

This keyword is related to the topics:

- problem

Specification

Alias: sample_refinement
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Required*(Choose One)*	**Group 1**	import	Sampling option
			adapt_import	Importance sampling option
			mm_adapt_import	Sampling option
	Optional		refinement_-samples	Specify the number of samples used to improve a probabilty estimate.

Description

Allow refinement of probability and generalized reliability results using importance sampling

The `probability_refinement` density reweighting must account for additional non-Gaussian cases. This option allows for refinement of probability and generalized reliability results using importance sampling.

import

- Keywords Area
- method
- stoch_collocation
- probability_refinement
- import

Sampling option

Specification

Alias: none

Argument(s): none

Description

`import` centers a sampling density at one of the initial LHS samples identified in the failure region. It then generates the importance samples, weights them by their probability of occurence given the original density, and calculates the required probability (CDF or CCDF level).

adapt_import

- Keywords Area
- method
- stoch_collocation
- probability_refinement
- adapt_import

Importance sampling option

Specification

Alias: none
Argument(s): none

Description

`adapt_import` centers a sampling density at one of the initial LHS samples identified in the failure region. It then generates the importance samples, weights them by their probability of occurence given the original density, and calculates the required probability (CDF or CCDF level). This continues iteratively until the failure probability estimate converges.

mm_adapt_import

- Keywords Area
- method
- stoch_collocation
- probability_refinement
- mm_adapt_import

Sampling option

Specification

Alias: none
Argument(s): none

Description

`mm_adapt_import` starts with all of the samples located in the failure region to build a multimodal sampling density. First, it uses a small number of samples around each of the initial samples in the failure region. Note that these samples are allocated to the different points based on their relative probabilities of occurrence: more probable points get more samples. This early part of the approach is done to search for "representative" points. Once these are located, the multimodal sampling density is set and then `mm_adapt_import` proceeds similarly to `adapt_import` (sample until convergence).

refinement_samples

- Keywords Area
- method
- stoch_collocation
- probability_refinement
- refinement_samples

Specify the number of samples used to improve a probabilty estimate.

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples used to improve a probabilty estimate. If using uni-modal sampling all samples are assigned to the sampling center. If using multi-modal sampling the samples are split between mutiple samples according to some internally computed weights.

export_points_file

- Keywords Area
- method
- stoch_collocation
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- stoch_collocation
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- stoch_collocation
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

fixed_seed

- Keywords Area
- method
- stoch_collocation
- fixed_seed

Reuse the same seed value for multiple sampling sets

Specification

Alias: none

Argument(s): none

Description

The `fixed_seed` flag is relevant if multiple sampling sets will be generated. This occurs when using advance methods (e.g., surrogate-based optimization, optimization under uncertainty).

Specifying this keyword results in the reuse of the same seed value for each of these multiple sampling sets, which can be important for reducing variability in the sampling results.

However, this behavior is not the default as the repetition of the same sampling pattern can result in a modeling weakness that an optimizer could potentially exploit (resulting in actual reliabilities that are lower than the estimated reliabilities). In either case (`fixed_seed` or not), the study is repeatable if the user specifies a `seed` and the study is random is the user omits a `seed` specification.

reliability_levels

- Keywords Area
- method
- stoch_collocation
- reliability_levels

Specify reliability levels at which the response values will be estimated

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_reliability_-levels	Specify which `reliability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF reliabilities by projecting out the prescribed number of sample standard deviations from the sample mean.

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_reliability_levels

- Keywords Area
- method
- stoch_collocation
- reliability_levels
- num_reliability_levels

Specify which `reliability_levels` correspond to which response

Specification

Alias: none

Argument(s): INTEGERLIST

Description

See parent page

response_levels

- Keywords Area
- method
- stoch_collocation
- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_-levels` correspond to which response
	Optional		compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_-response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- stoch_collocation
- response_levels
- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- stoch_collocation
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	probabilities	Compute probabilities
			reliabilities	Compute reliabilities
			gen_reliabilities	Compute generalized reliabilities
	Optional		system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- stoch_collocation
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

reliabilities

- Keywords Area
- method
- stoch_collocation
- response_levels
- compute

- reliabilities

Compute reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- stoch_collocation
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- stoch_collocation
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	series	?????
			parallel	?????

series

- Keywords Area
- method
- stoch_collocation
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): none

parallel

- Keywords Area
- method
- stoch_collocation
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

distribution

- Keywords Area
- method
- stoch_collocation
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional	**Description of Group**	**Dakota Keyword**	**Dakota Keyword Description**
	Required*(Choose One)*	**Group 1**	cumulative	Choose cumulative functions

			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.
This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- stoch_collocation
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- stoch_collocation
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- stoch_collocation
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- stoch_collocation
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- stoch_collocation
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- stoch_collocation
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none

Argument(s): INTEGERLIST

Description

See parent page

rng

- Keywords Area
- method
- stoch_collocation
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional Required*(Choose One)*	Description of Group Group 1	Dakota Keyword mt19937	Dakota Keyword Description Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- stoch_collocation
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: `http://en.wikipedia.org/wiki/Mersenne_twister`

rnum2

- Keywords Area
- method
- stoch_collocation
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

samples

- Keywords Area
- method
- stoch_collocation
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- stoch_collocation
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- stoch_collocation
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.53 sampling

- Keywords Area
- method
- sampling

Uncertainty quantification based on random sampling of the variables

Topics

This keyword is related to the topics:

- uncertainty_quantification
- sampling

Specification

Alias: nond_sampling

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		sample_type	Choose the sampling algorithm
	Optional		variance_based_-decomp	Global sensitivity analyis using variance-based decomposition
	Optional		fixed_seed	Reuse the same seed value for multiple sampling sets
	Optional		reliability_levels	Specify reliability levels at which the response values will be estimated
	Optional		response_levels	Specify the response values at which to estimate relevant statistics
	Optional		distribution	Select between cumulative or complementary-cumulative functions

	Optional	probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional	gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional	rng	Choose a random number generator
	Optional	samples	Specify the number of samples
	Optional	seed	Specify the seed of the random number generator
	Optional	model_pointer	Specify which model block will be used by a method

Description

This method draws samples from the specified uncertain variable probability distributions and propagates them through the model to obtain statistics on the distributions of the responses.

Method Behavior

The following keywords change how the samples are selected:

- sample_type
- fixed_seed
- rng
- samples
- seed
- variance_based_decomp

Method Outputs

As a default, Dakota provides correlation analyses when running LHS. Correlation tables are printed with the simple, partial, and rank correlations between inputs and outputs. These can be useful to get a quick sense of how correlated the inputs are to each other, and how correlated various outputs are to inputs. `variance_based_-decomp` is used to request more sensitivity information, with additional cost.

Additional statistics can be computed from the samples using the following keywords:

- `response_levels`
- `reliability_levels`
- `probability_levels`

- gen_reliability_levels

response_levels computes statistics at the specified response value. The other three allow the specification of the statistic value, and will estimate the corresponding response value.

distribution is used to specify whether the statistic values are from cumulative or complementary cumulative functions.

Examples

```
# tested on Dakota 6.0 on 140501

environment
  tabular_graphics_data
    tabular_graphics_file = 'Sampling_basic.dat'

method
  sampling
    sample_type lhs
    samples = 20

model
  single

variables
  active uncertain
  uniform_uncertain = 2
    descriptors  =   'input1'     'input2'
    lower_bounds =  -2.0     -2.0
    upper_bounds =   2.0      2.0
  continuous_state = 1
    descriptors =   'constant1'
    initial_state = 100

interface
  analysis_drivers 'text_book'
    fork

responses
  response_functions = 1
  no_gradients
  no_hessians
```

This example illustrates a basic sampling Dakota input file.

- LHS is used instead of purely random sampling.
- The default random number generator is used.
- Without a seed specified, this will not be reproducable
- In the variables block, two types of variables are used
- Only the uncertain variables are varied, this is the default behavior, and is also specified by the active keyword, w/ the uncertain option

See Also

These keywords may also be of interest:

- active
- incremental_lhs

FAQ

Q: Do I need to keep the LHS∗ and S4 files? **A:** No

sample_type

- Keywords Area
- method
- sampling
- sample_type

Choose the sampling algorithm

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	random	Purely random samples
			lhs	Use LHS method
			incremental_lhs	Augment an existing LHS study
			incremental_-random	Augment an existing random sampling study

Description

There are only two types of sampling:

- Monte Carlo (pure random)
- Latin Hypercube Sampling

However, some methods allow an existing study to be augmented with additional or incremental sampling. This is used to get better estimates of mean, variance, and percentiles. The incremental study must use the same type of sampling as the existing study.

random

- Keywords Area
- method
- sampling
- sample_type
- random

Purely random samples

Specification

Alias: none
Argument(s): none

Description

Purely random samples.

lhs

- Keywords Area
- method
- sampling
- sample_type
- lhs

Use LHS method

Specification

Alias: none
Argument(s): none

Description

Use Latin Hypercube Sampling

incremental_lhs

- Keywords Area
- method
- sampling
- sample_type
- incremental_lhs

Augment an existing LHS study

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		previous_samples	Number of samples in a previous sampling study

Description

`incremental_lhs` will augment an existing LHS sampling study with more samples to get better estimates of mean, variance, and percentiles.

The number of samples in the second set MUST currently be 2 times the number of previous samples, although incremental sampling based on any power of two may be supported in future releases.

To use this `sample_type`, a previous LHS (or incremental LHS) sampling study with sample size N must have already been performed, and **the dakota restart file must be available from this previous study.** The variables and responses specifications must be the same in both studies.

See the examples below and the Usage and Restarting Dakota Studies pages.

Examples

For example, say a user performs an initial study using `lhs` as the `sample_type`, and generates 10 samples.

One way to ensure the restart file is saved is to specify a non-default name, via a command line option:

```
dakota -i LHS_10.in -w LHS_10.rst
```

which uses the input file:

```
# LHS_10.in

environment
  tabular_graphics_data
    tabular_graphics_file = 'lhs10.dat'

method
  sampling
    sample_type lhs
    samples = 10

model
  single

variables
  uniform_uncertain = 2
    descriptors  =   'input1'     'input2'
    lower_bounds =  -2.0     -2.0
    upper_bounds =   2.0      2.0

interface
  analysis_drivers 'text_book'
    fork

responses
  response_functions = 1
```

```
  no_gradients
  no_hessians
```

and the restart file is written to LHS_10.rst.

Then an incremental LHS study can be run with:

```
dakota -i LHS_20.in -r LHS_10.rst -w LHS_20.rst
```

where `LHS_20.in` is shown below, and LHS_10.rst is the restart file containing the results of the previous LHS study.

```
# LHS_20.in

environment
  tabular_graphics_data
    tabular_graphics_file = 'lhs_incremental_20.dat'

method
  sampling
    sample_type incremental_lhs
    samples = 20
    previous_samples = 10

model
  single

variables
  uniform_uncertain = 2
    descriptors  =   'input1'     'input2'
    lower_bounds =  -2.0     -2.0
    upper_bounds =   2.0      2.0

interface
  analysis_drivers 'text_book'
    fork

responses
  response_functions = 1
  no_gradients
  no_hessians
```

The user will get 10 new LHS samples which maintain both the correlation and stratification of the original LHS sample. The new samples will be combined with the original samples to generate a combined sample of size 20.

This is clearly seen by comparing the two tabular graphics data files.

previous_samples

- Keywords Area
- method
- sampling
- sample_type
- incremental_lhs
- previous_samples

Number of samples in a previous sampling study

Specification

Alias: none
Argument(s): INTEGER

Description

See the parent page for information.

incremental_random

- Keywords Area
- method
- sampling
- sample_type
- incremental_random

Augment an existing random sampling study

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		previous_samples	Number of samples in a previous sampling study

Description

`incremental_random` will augment an existing random sampling study with more samples to get better estimates of mean, variance, and percentiles.

To use this `sample_type`, a previous random sampling study with sample size N must have already been performed, and the dakota restart file must be available from this previous study. The variables and responses specifications must be the same in both studies.

The syntax for running the second sample set is:

```
dakota -i input2.in -r dakota.rst
```

where `input2.in` is the file which specifies incremental sampling and dakota.rst is the restart file containing the results of the previous study.

The number of samples in the second set MUST currently be 2 times the number of previous samples, although incremental sampling based on any power of two may be supported in future releases.

Examples

For example, say a user performs an initial study using `random` as the `sample_type`, and generates 50 samples. If the user creates a new input file where `samples` is now specified to be 100, the `sample_type` is defined to be `incremental_random`, and `previous_samples` is specified to be 50, the user will get 50 new random samples. The N new samples will be combined with the N previous samples to generate a combined sample of size 2N.

previous_samples

- Keywords Area
- method
- sampling
- sample_type
- incremental_random
- previous_samples

Number of samples in a previous sampling study

Specification

Alias: none
Argument(s): INTEGER

Description

See the parent page for information.

variance_based_decomp

- Keywords Area
- method
- sampling
- variance_based_decomp

Global sensitivity analyis using variance-based decomposition

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Optional	drop_tolerance	Drop values lower than this tolerance

Description

Dakota can calculate sensitivity indices through variance based decomposition using the keyword `variance_based_decomp`. These indicate how important the uncertainty in each input variable is in contributing to the output variance.

`variance_based_decomp` is extremely computationally intensive since replicated sets of sample values are evaluated. If the user specified a number of samples, N, and a number of nondeterministic variables, M, variance-based decomposition requires the evaluation of N∗(M+2) samples. **Note that specifying this keyword will increase the number of function evaluations above the number requested with the `samples` keyword** .

Because of the computational cost, `variance_based_decomp` is turned off as a default.

To obtain sensitivity indices that are reasonably accurate, we recommend that N, the number of samples, be at least one hundred and preferably several hundred or thousands.

Theory

In this context, we take sensitivity analysis to be global, not local as when calculating derivatives of output variables with respect to input variables. Our definition is similar to that of [[74] "Saltelli et al., 2004"]: "The study of how uncertainty in the output of a model can be apportioned to different sources of uncertainty in the model input."

Variance based decomposition is a way of using sets of samples to understand how the variance of the output behaves, with respect to each input variable. A larger value of the sensitivity index, S_i, means that the uncertainty in the input variable i has a larger effect on the variance of the output. More details on the calculations and interpretation of the sensitivity indices can be found in [[74] "Saltelli et al., 2004"] and [[88] "Weirs et al., 2010"].

drop_tolerance

- Keywords Area
- method
- sampling
- variance_based_decomp
- drop_tolerance

Drop values lower than this tolerance

Specification

Alias: none

Argument(s): REAL

Description

Drop values lower than this tolerance. For PCE which outputs main, interaction, and total effects by default, the univariate_effects option has been added to allow suppression of the interaction effects since the output volume of these results can be prohibitive for high dimensional problems. Similar to suppression of these interactions is the covariance control, which can be selected to be diagonal_covariance or full_covariance, with the former

supporting suppression of the off-diagonal covariance terms (to save compute and memory resources and reduce output volume).

fixed_seed

- Keywords Area
- method
- sampling
- fixed_seed

Reuse the same seed value for multiple sampling sets

Specification

Alias: none
Argument(s): none

Description

The `fixed_seed` flag is relevant if multiple sampling sets will be generated. This occurs when using advance methods (e.g., surrogate-based optimization, optimization under uncertainty).

Specifying this keyword results in the reuse of the same seed value for each of these multiple sampling sets, which can be important for reducing variability in the sampling results.

However, this behavior is not the default as the repetition of the same sampling pattern can result in a modeling weakness that an optimizer could potentially exploit (resulting in actual reliabilities that are lower than the estimated reliabilities). In either case (`fixed_seed` or not), the study is repeatable if the user specifies a `seed` and the study is random is the user omits a `seed` specification.

reliability_levels

- Keywords Area
- method
- sampling
- reliability_levels

Specify reliability levels at which the response values will be estimated

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Optional	num_reliability_-levels	Specify which `reliability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF reliabilities by projecting out the prescribed number of sample standard deviations from the sample mean.

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_reliability_levels

- Keywords Area
- method
- sampling
- reliability_levels
- num_reliability_levels

Specify which `reliability_levels` correspond to which response

Specification

Alias: none

Argument(s): INTEGERLIST

Description

See parent page

response_levels

- Keywords Area
- method
- sampling
- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_levels` correspond to which response
	Optional		compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- sampling
- response_levels
- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- sampling
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none
Argument(s): none

	Required/-Optional	**Description of Group**	**Dakota Keyword**	**Dakota Keyword Description**
	Required*(Choose One)*	**Group 1**	probabilities	Compute probabilities
			reliabilities	Compute reliabilities
			gen_reliabilities	Compute generalized reliabilities

	Optional	system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- sampling
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

reliabilities

- Keywords Area
- method
- sampling
- response_levels
- compute
- reliabilities

Compute reliabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- sampling
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- sampling
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	series	?????
			parallel	?????

series

- Keywords Area
- method
- sampling
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

parallel

- Keywords Area
- method
- sampling
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

distribution

- Keywords Area
- method
- sampling
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	cumulative	Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.

This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- sampling
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- sampling
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- sampling
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- sampling
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- sampling
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- sampling
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

rng

- Keywords Area
- method
- sampling
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** mt19937	**Dakota Keyword Description** Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- sampling
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: http://en.wikipedia.org/wiki/Mersenne_twister

rnum2

- Keywords Area
- method
- sampling
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

samples

- Keywords Area
- method
- sampling
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- sampling
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- sampling
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.54 importance_sampling

- Keywords Area
- method
- importance_sampling

Importance sampling

Topics

This keyword is related to the topics:

- uncertainty_quantification
- aleatory_uncertainty_quantification_methods
- sampling

Specification

Alias: nond_importance_sampling
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	import	Sampling option
			adapt_import	Importance sampling option
			mm_adapt_import	Sampling option
	Optional		refinement_-samples	Specify the number of samples used to improve a probabilty estimate.
	Optional		response_levels	Specify the response values at which to estimate relevant statistics
	Optional		distribution	Select between cumulative or complementary-cumulative functions
	Optional		probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional		gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional		rng	Choose a random number generator
	Optional		samples	Specify the number of samples
	Optional		seed	Specify the seed of the random number generator
	Optional		model_pointer	Specify which model block will be used by a method

Description

The `importance_sampling` method is based on ideas in reliability modeling.

An initial Latin Hypercube sampling is performed to generate an initial set of samples. These initial samples are augmented with samples from an importance density as follows:

- The variables are transformed to standard normal space.
- In the transformed space, the importance density is a set of normal densities centered around points which are in the failure region.
- Note that this is similar in spirit to the reliability methods, in which importance sampling is centered around a Most Probable Point (MPP).
- In the case of the LHS samples, the importance sampling density will simply by a mixture of normal distributions centered around points in the failure region.

Options

Choose one of the importance sampling options:

- `import`
- `adapt_import`
- `mm_adapt_import`

The options for importance sampling are as follows: `import` centers a sampling density at one of the initial LHS samples identified in the failure region. It then generates the importance samples, weights them by their probability of occurence given the original density, and calculates the required probability (CDF or CCDF level). `adapt_import` is the same as `import` but is performed iteratively until the failure probability estimate converges. `mm_adapt_import` starts with all of the samples located in the failure region to build a multimodal sampling density. First, it uses a small number of samples around each of the initial samples in the failure region. Note that these samples are allocated to the different points based on their relative probabilities of occurrence: more probable points get more samples. This early part of the approach is done to search for ”representative” points. Once these are located, the multimodal sampling density is set and then `mm_adapt_import` proceeds similarly to `adapt_import` (sample until convergence).

Theory

Importance sampling is a method that allows one to estimate statistical quantities such as failure probabilities (e.g. the probability that a response quantity will exceed a threshold or fall below a threshold value) in a way that is more efficient than Monte Carlo sampling. The core idea in importance sampling is that one generates samples that preferentially samples important regions in the space (e.g. in or near the failure region or user-defined region of interest), and then appropriately weights the samples to obtain an unbiased estimate of the failure probability [[77] ”Srinivasan, 2002”]. In importance sampling, the samples are generated from a density which is called the importance density: it is not the original probability density of the input distributions. The importance density should be centered near the failure region of interest. For black-box simulations such as those commonly interfaced with Dakota, it is difficult to specify the importance density a priori: the user often does not know where the failure region lies, especially in a high-dimensional space.[[79] ”Swiler and West, 2010”]. We have developed two importance sampling approaches which do not rely on the user explicitly specifying an importance density.

See Also

These keywords may also be of interest:

- adaptive_sampling
- gpais
- local_reliability
- global_reliability
- sampling
- polynomial_chaos
- stoch_collocation

import

- Keywords Area
- method
- importance_sampling
- import

Sampling option

Specification

Alias: none
Argument(s): none

Description

`import` centers a sampling density at one of the initial LHS samples identified in the failure region. It then generates the importance samples, weights them by their probability of occurence given the original density, and calculates the required probability (CDF or CCDF level).

adapt_import

- Keywords Area
- method
- importance_sampling
- adapt_import

Importance sampling option

Specification

Alias: none
Argument(s): none

Description

adapt_import centers a sampling density at one of the initial LHS samples identified in the failure region. It then generates the importance samples, weights them by their probability of occurence given the original density, and calculates the required probability (CDF or CCDF level). This continues iteratively until the failure probability estimate converges.

mm_adapt_import

- Keywords Area
- method
- importance_sampling
- mm_adapt_import

Sampling option

Specification

Alias: none
Argument(s): none

Description

mm_adapt_import starts with all of the samples located in the failure region to build a multimodal sampling density. First, it uses a small number of samples around each of the initial samples in the failure region. Note that these samples are allocated to the different points based on their relative probabilities of occurrence: more probable points get more samples. This early part of the approach is done to search for "representative" points. Once these are located, the multimodal sampling density is set and then mm_adapt_import proceeds similarly to adapt_import (sample until convergence).

refinement_samples

- Keywords Area
- method
- importance_sampling
- refinement_samples

Specify the number of samples used to improve a probabilty estimate.

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples used to improve a probabilty estimate. If using uni-modal sampling all samples are assigned to the sampling center. If using multi-modal sampling the samples are split between mutiple samples according to some internally computed weights.

response_levels

- Keywords Area
- method
- importance_sampling
- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_-levels` correspond to which response
	Optional		compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- importance_sampling
- response_levels
- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none

Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- importance_sampling
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	probabilities	Compute probabilities
			gen_reliabilities	Compute generalized reliabilities
	Optional		system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- importance_sampling
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- importance_sampling
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- importance_sampling
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	series	?????
			parallel	?????

series

- Keywords Area
- method
- importance_sampling
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

parallel

- Keywords Area
- method
- importance_sampling
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

distribution

- Keywords Area
- method
- importance_sampling
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	cumulative	Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.

This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- importance_sampling
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- importance_sampling
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- importance_sampling
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- importance_sampling
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- importance_sampling
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- importance_sampling
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

rng

- Keywords Area
- method
- importance_sampling
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** mt19937	**Dakota Keyword Description** Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- importance_sampling
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: http://en.wikipedia.org/wiki/Mersenne_twister

rnum2

- Keywords Area
- method
- importance_sampling
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

samples

- Keywords Area
- method
- importance_sampling
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- importance_sampling
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- importance_sampling
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.55 gpais

- Keywords Area
- method
- gpais

Gaussian Process Adaptive Importance Sampling

Topics

This keyword is related to the topics:

- uncertainty_quantification

Specification

Alias: gaussian_process_adaptive_importance_sampling
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		emulator_samples	Number of data points used to train the surrogate model or emulator
	Optional		import_points_file	Specify file containing points (variable and response values)

	Optional	export_points_file	Specify file to write out the responses from the surrogate
	Optional	response_levels	Specify the response values at which to estimate relevant statistics
	Optional	distribution	Select between cumulative or complementary-cumulative functions
	Optional	probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional	gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional	rng	Choose a random number generator
	Optional	samples	Specify the number of samples
	Optional	seed	Specify the seed of the random number generator
	Optional	model_pointer	Specify which model block will be used by a method

Description

`gpais` is recommended for problems that have a relatively small number of input variables (e.g. less than 10-20). This method, Gaussian Process Adaptive Importance Sampling, is outlined in the paper [[**?**] "Dalbey and Swiler, 2012"].

This method starts with an initial set of LHS samples and adds samples one at a time, with the goal of adaptively improving the estimate of the ideal importance density during the process. The approach uses a mixture of component densities. An iterative process is used to construct the sequence of improving component densities. At each iteration, a Gaussian process (GP) surrogate is used to help identify areas in the space where failure is likely to occur. The GPs are not used to directly calculate the failure probability; they are only used to approximate the importance density. Thus, the Gaussian process adaptive importance sampling algorithm overcomes limitations involving using a potentially inaccurate surrogate model directly in importance sampling calculations.

See Also

These keywords may also be of interest:

- adaptive_sampling
- local_reliability
- global_reliability
- sampling
- importance_sampling
- polynomial_chaos
- stoch_collocation

emulator_samples

- Keywords Area
- method
- gpais
- emulator_samples

Number of data points used to train the surrogate model or emulator

Specification

Alias: none
Argument(s): INTEGER

Description

This keyword refers to the number of build points or training points used to construct a Gaussian process emulator. If the user specifies a number of `emulator_samples` that is less than the minimum number of points required to build the GP surrogate, Dakota will augment the samples to obtain the minimum required.

import_points_file

- Keywords Area
- method
- gpais
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- gpais
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- gpais
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- gpais
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none
Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- gpais
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- gpais
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

response_levels

- Keywords Area
- method
- gpais
- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_levels` correspond to which response
	Optional		compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- gpais
- response_levels
- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- gpais
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	probabilities	Compute probabilities
			gen_reliabilities	Compute generalized reliabilities
	Optional		system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- gpais
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- gpais
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- gpais
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	series	?????
			parallel	?????

series

- Keywords Area
- method
- gpais
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

parallel

- Keywords Area
- method
- gpais
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

distribution

- Keywords Area
- method
- gpais
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	cumulative	Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.
This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- gpais
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- gpais
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- gpais
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- gpais
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- gpais
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- gpais
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none

Argument(s): INTEGERLIST

Description

See parent page

rng

- Keywords Area
- method
- gpais
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional Required*(Choose One)*	Description of Group Group 1	Dakota Keyword mt19937	Dakota Keyword Description Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- gpais
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: `http://en.wikipedia.org/wiki/Mersenne_twister`

rnum2

- Keywords Area
- method
- gpais
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

samples

- Keywords Area
- method
- gpais
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- gpais
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- gpais
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.56 adaptive_sampling

- Keywords Area
- method
- adaptive_sampling

Build a GP surrogate and refine it adaptively

Topics

This keyword is related to the topics:

- uncertainty_quantification

Specification

Alias: nond_adaptive_sampling
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		emulator_samples	Number of data points used to train the surrogate model or emulator
	Optional		fitness_metric	Specify the `fitness_metric` used to select the next point
	Optional		batch_selection	How to select new points
	Optional		batch_size	The number of points to add in each batch.
	Optional		import_points_file	Specify file containing points (variable and response values)
	Optional		export_points_file	Specify file to write out the responses from the surrogate

	Optional	response_levels	Specify the response values at which to estimate relevant statistics
	Optional	misc_options	
	Optional	distribution	Select between cumulative or complementary-cumulative functions
	Optional	probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional	gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional	rng	Choose a random number generator
	Optional	samples	Specify the number of samples
	Optional	seed	Specify the seed of the random number generator
	Optional	model_pointer	Specify which model block will be used by a method

Description

The goal in performing adaptive sampling is to construct a surrogate model that can be used as an accurate predictor to some expensive simulation, thus it is to one's advantage to build a surrogate that minimizes the error over the entire domain of interest using as little data as possible from the expensive simulation. The adaptive part alludes to the fact that the surrogate will be refined by focusing samples of the expensive simulation on particular areas of interest rather than rely on random selection or standard space-filling techniques.

At a high-level, the adaptive sampling pipeline is a four-step process:

- Evaluate the expensive simulation (referred to as the true model) at initial sample point

  1. Fit a surrogate model
  2. Create a candidate set and score based on information from surrogate
  3. Select a candidate point to evaluate the true model
  4. Loop until done

In terms of the Dakota implementation, the adaptive sampling method currently uses Latin Hypercube sampling (LHS) to generate the initial points in Step 1 above. For Step 2, we use a Gaussian process model.

The default behavior is to add one point at a time. At each iteration (e.g. each loop of Steps 2-4 above), a Latin Hypercube sample is generated (a new one, different from the initial sample) and the surrogate model is evaluated at this points. These are the candidate points that are then evaluated according to the fitness metric. The number of candidates used in practice should be high enough to fill most of the input domain: we recommend at least hundreds of points for a low-dimensional problem. All of the candidates (samples on the emulator) are given a score and then the highest-scoring candidate is selected to be evaluated on the true model.

The adaptive sampling method also can generate batches of points to add at a time using the `batch_selection` and `batch_size` keywords.

See Also

These keywords may also be of interest:

- gpais
- local_reliability
- global_reliability
- sampling
- importance_sampling
- polynomial_chaos
- stoch_collocation

emulator_samples

- Keywords Area
- method
- adaptive_sampling
- emulator_samples

Number of data points used to train the surrogate model or emulator

Specification

Alias: none
Argument(s): INTEGER

Description

This keyword refers to the number of build points or training points used to construct a Gaussian process emulator. If the user specifies a number of `emulator_samples` that is less than the minimum number of points required to build the GP surrogate, Dakota will augment the samples to obtain the minimum required.

fitness_metric

- Keywords Area
- method
- adaptive_sampling
- fitness_metric

Specify the `fitness_metric` used to select the next point

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	predicted_variance	Pick points with highest variance
			distance	Space filling metric
			gradient	Fill the range space of the surrogate

Description

The user can specify the `fitness_metric` used to select the next point (or points) to evaluate and add to the set. The fitness metrics used for scoring candidate points include:

- `predicted_variance`
- `distance`
- `gradient`

predicted_variance

- Keywords Area
- method
- adaptive_sampling
- fitness_metric
- predicted_variance

Pick points with highest variance

Specification

Alias: none

Argument(s): none

Description

The predicted variance metric uses the predicted variance of the Gaussian process surrogate as the score of a candidate point. Thus, the adaptively chosen points will be in areas of highest uncertainty according to the Gaussian process model.

distance

- Keywords Area
- method
- adaptive_sampling
- fitness_metric
- distance

Space filling metric

Specification

Alias: none
Argument(s): none

Description

The distance metric calculates the Euclidean distance in domain space between the candidate and its nearest neighbor in the set of points already evaluated on the true model. Therefore, the most undersampled area of the domain will always be selected. Note that this is a space-filling metric.

gradient

- Keywords Area
- method
- adaptive_sampling
- fitness_metric
- gradient

Fill the range space of the surrogate

Specification

Alias: none
Argument(s): none

Description

The gradient metric calculates the score as the absolute value of the difference in range space (the outputs) of the two points. The output values used are predicted from the surrogate model. This method attempts to evenly fill the range space of the surrogate.

batch_selection

- Keywords Area
- method
- adaptive_sampling
- batch_selection

How to select new points

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	naive	Take the highest scoring candidates
			distance_penalty	Add a penalty to spread out the points in the batch
			topology	In this selection strategy, we use information about the topology of the space from the Morse-Smale complex to identify next points to select.
			constant_liar	Use information from the existing surrogate model to predict what the surrogate upgrade will be with new points.

Description

With batch or multi-point selection, the true model can be evaluated in parallel and thus increase throughput before refitting our surrogate model. This proposes a new challenge as the problem of choosing a single point and choosing multiple points off a surrogate are fundamentally different. Selecting the `n` best scoring candidates is more than likely to generate a set of points clustered in one area which will not be conducive to adapting the surrogate.

We have implemented several strategies for batch selection of points. These are described in the User's manual and are the subject of active research.

The `batch_selection` strategies include:

1. `naive:`

2. `distance_penalty`
3. `constant_liar`
4. `topology`

naive

- Keywords Area
- method
- adaptive_sampling
- batch_selection
- naive

Take the highest scoring candidates

Specification

Alias: none
Argument(s): none

Description

This strategy will select the `n` highest scoring candidates regardless of their position. This tends to group an entire round of points in the same area.

distance_penalty

- Keywords Area
- method
- adaptive_sampling
- batch_selection
- distance_penalty

Add a penalty to spread out the points in the batch

Specification

Alias: none
Argument(s): none

Description

In this strategy, the highest scoring candidate is selected and then all remaining candidates are re-scored with a distance penalization factor added in to the score.

topology

- Keywords Area
- method
- adaptive_sampling
- batch_selection
- topology

In this selection strategy, we use information about the topology of the space from the Morse-Smale complex to identify next points to select.

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): none

Description

In this strategy we look at the topology of the scoring function and select the `n` highest maxima in the topology. To determine local maxima, we construct the approximate Morse-Smale complex. This strategy does require the user to have the Morse-Smale package.

constant_liar

- Keywords Area
- method
- adaptive_sampling
- batch_selection
- constant_liar

Use information from the existing surrogate model to predict what the surrogate upgrade will be with new points.

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): none

Description

The strategy first selects the highest scoring candidate, and then refits the surrogate using a "lie" value at the point selected and repeats until `n` points have been selected whereupon the lie values are removed from the surrogate and the selected points are evaluated on the true model and the surrogate is refit with these values.

batch_size

- Keywords Area
- method
- adaptive_sampling
- batch_size

The number of points to add in each batch.

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGER

Description

The number of points to add in each batch.

import_points_file

- Keywords Area
- method
- adaptive_sampling
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- adaptive_sampling
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- adaptive_sampling
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- adaptive_sampling
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none
Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- adaptive_sampling
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- adaptive_sampling
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

response_levels

- Keywords Area
- method
- adaptive_sampling
- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_levels` correspond to which response
	Optional		compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- adaptive_sampling
- response_levels
- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- adaptive_sampling
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	probabilities	Compute probabilities
			gen_reliabilities	Compute generalized reliabilities
	Optional		system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- adaptive_sampling
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- adaptive_sampling
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- adaptive_sampling
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	series	?????
			parallel	?????

series

- Keywords Area
- method
- adaptive_sampling
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

parallel

- Keywords Area
- method
- adaptive_sampling
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

misc_options

- Keywords Area
- method
- adaptive_sampling
- misc_options

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): STRINGLIST

distribution

- Keywords Area
- method
- adaptive_sampling
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** cumulative	**Dakota Keyword Description** Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.
This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- adaptive_sampling
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- adaptive_sampling
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- adaptive_sampling
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- adaptive_sampling
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- adaptive_sampling
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- adaptive_sampling
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

rng

- Keywords Area
- method
- adaptive_sampling
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** mt19937	**Dakota Keyword Description** Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- adaptive_sampling
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: http://en.wikipedia.org/wiki/Mersenne_twister

rnum2

- Keywords Area
- method
- adaptive_sampling
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

samples

- Keywords Area
- method
- adaptive_sampling
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- adaptive_sampling
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- adaptive_sampling
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.57 pof_darts

- Keywords Area
- method
- pof_darts

Probability-of-Failure (POF) darts is a novel method for estimating the probability of failure based on random sphere-packing.

Topics

This keyword is related to the topics:

- uncertainty_quantification

Specification

Alias: nond_pof_darts
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		response_levels	Specify the response values at which to estimate relevant statistics
	Optional		distribution	Select between cumulative or complementary-cumulative functions

	Optional	probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional	gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional	rng	Choose a random number generator
	Optional	samples	Specify the number of samples
	Optional	seed	Specify the seed of the random number generator
	Optional	model_pointer	Specify which model block will be used by a method

Description

`pof_darts` is a novel method for estimating the probability of failure based on random sphere-packing. Random spheres are sampled from the domain with the constraint that each new sphere center has to be outside prior disks. The radius of each sphere is chosen such that the entire sphere lie either in the failure or the non-failure region. This radius depends of the function evaluation at the disk center, the failure threshold and an estimate of the function gradient at the disk center. We utilize a Gaussian process surrogate for evaluating the gradient and hence only one function evaluation is required for each sphere.

After exhausting the sampling budget specified by `samples`, which is the number of spheres per failure threshold, the domain is decomposed into two regions. These regions correspond to failure and non-failure, each represented by the union of the spheres of each type. The volume of the union of failure spheres gives a lower bound on the required estimate of the probability of failure, while the volume of the union of the non-failure spheres subtracted from the volume of the domain gives an upper estimate. We currently report the average of both estimates.

`pof_darts` handles multiple response functions and allows each to have multiple failure thresholds. For each failure threshold, `pof_darts` will insert a number of spheres specified by the user-input parameter `samples`. However, estimating the probability of failure for each failure threshold would utilize the total number of disks sampled for all failure thresholds. For each failure threshold, the sphere radii changes to generate the right spatial decomposition.

response_levels

- Keywords Area
- method
- pof_darts

- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_levels` correspond to which response
	Optional		compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- pof_darts
- response_levels
- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- pof_darts
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	probabilities	Compute probabilities
			gen_reliabilities	Compute generalized reliabilities
	Optional		system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- pof_darts
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- pof_darts
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- pof_darts
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	series	?????
			parallel	?????

series

- Keywords Area
- method
- pof_darts
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

parallel

- Keywords Area
- method
- pof_darts
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

distribution

- Keywords Area
- method
- pof_darts
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	cumulative	Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.
This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- pof_darts
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- pof_darts
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- pof_darts
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- pof_darts
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- pof_darts
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- pof_darts
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none

Argument(s): INTEGERLIST

Description

See parent page

rng

- Keywords Area
- method
- pof_darts
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional Required*(Choose One)*	Description of Group Group 1	Dakota Keyword mt19937	Dakota Keyword Description Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- pof_darts
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: `http://en.wikipedia.org/wiki/Mersenne_twister`

rnum2

- Keywords Area
- method
- pof_darts
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

samples

- Keywords Area
- method
- pof_darts
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- pof_darts
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- pof_darts
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.58 efficient_subspace

- Keywords Area
- method
- efficient_subspace

This is not mentioned in the reference manual

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: nond_efficient_subspace

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		emulator_samples	Number of data points used to train the surrogate model or emulator
	Optional		batch_size	The number of points to add in each batch.
	Optional		distribution	Select between cumulative or complementary-cumulative functions
	Optional		probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional		gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value

	Optional	rng	Choose a random number generator
	Optional	samples	Specify the number of samples
	Optional	seed	Specify the seed of the random number generator
	Optional	model_pointer	Specify which model block will be used by a method

emulator_samples

- Keywords Area
- method
- efficient_subspace
- emulator_samples

Number of data points used to train the surrogate model or emulator

Specification

Alias: none
Argument(s): INTEGER

Description

This keyword refers to the number of build points or training points used to construct a Gaussian process emulator. If the user specifies a number of `emulator_samples` that is less than the minimum number of points required to build the GP surrogate, Dakota will augment the samples to obtain the minimum required.

batch_size

- Keywords Area
- method
- efficient_subspace
- batch_size

The number of points to add in each batch.

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGER

Description

The number of points to add in each batch.

distribution

- Keywords Area
- method
- efficient_subspace
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** cumulative	**Dakota Keyword Description** Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.
This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- efficient_subspace
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- efficient_subspace
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- efficient_subspace
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- efficient_subspace
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- efficient_subspace
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- efficient_subspace
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

rng

- Keywords Area
- method
- efficient_subspace
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** mt19937	**Dakota Keyword Description** Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- efficient_subspace
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: http://en.wikipedia.org/wiki/Mersenne_twister

rnum2

- Keywords Area
- method
- efficient_subspace
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

samples

- Keywords Area
- method
- efficient_subspace
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- efficient_subspace
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- efficient_subspace
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.59 global_evidence

- Keywords Area
- method
- global_evidence

Evidence theory with evidence measures computed with global optimization methods

Topics

This keyword is related to the topics:

- problem
- epistemic_uncertainty_quantification_methods
- evidence_theory

Specification

Alias: nond_global_evidence
Argument(s): none

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** sbo	**Dakota Keyword Description** Use the surrogate based optimization method
			ego	Use the Efficient Global Optimization method

			ea	Use an evolutionary algorithm
			lhs	Use LHS method
	Optional		response_levels	Specify the response values at which to estimate relevant statistics
	Optional		distribution	Select between cumulative or complementary-cumulative functions
	Optional		probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional		gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional		rng	Choose a random number generator
	Optional		samples	Specify the number of samples
	Optional		seed	Specify the seed of the random number generator
	Optional		model_pointer	Specify which model block will be used by a method

Description

See the topic page evidence_theory for important background information and usage notes

`global_evidence` allows the user to specify several global approaches for calculating the belief and plausibility functions:

- `lhs` - note: this takes the minimum and maximum of the samples as the bounds per "interval cell combination."
- `ego`
- `sbo`
- `ea`

Note that to calculate the plausibility and belief cumulative distribution functions, one has to look at all combinations of intervals for the uncertain variables. In terms of implementation, if one is using LHS sampling as outlined above, this method creates a large sample over the response surface, then examines each cell to determine the minimum and maximum sample values within each cell. To do this, one needs to set the number of samples relatively high: the default is 10,000 and we recommend at least that number. If the model you are running is a simulation that is computationally quite expensive, we recommend that you set up a surrogate model within the Dakota input file so that `global_evidence` performs its sampling and calculations on the surrogate and not on the original model. If one uses optimization methods instead to find the minimum and maximum sample values within each cell, this can also be computationally expensive.

Theory

The basic idea is that one specifies an ”evidence structure” on uncertain inputs and propagates that to obtain belief and plausibility functions on the response functions. The inputs are defined by sets of intervals and Basic Probability Assignments (BPAs). Evidence propagation is computationally expensive, since the minimum and maximum function value must be calculated for each ”interval cell combination.” These bounds are aggregated into belief and plausibility.

See Also

These keywords may also be of interest:

- global_interval_est
- local_evidence
- local_interval_est

sbo

- Keywords Area
- method
- global_evidence
- sbo

Use the surrogate based optimization method

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		gaussian_process	Use a Gaussian Process surrogate

	Optional	use_derivatives	Use derivative data to construct surrogate models
	Optional	import_points_file	Specify file containing points (variable and response values)
	Optional	export_points_file	Specify file to write out the responses from the surrogate

Description

A surrogate-based optimization method will be used. The surrogate employed in `sbo` is a Gaussian process surrogate.

The main difference between `ego` and the `sbo` approach is the objective function being optimized. `ego` relies on an expected improvement function, while in `sbo`, the optimization proceeds using an evolutionary algorithm (coliny_ea) on the Gaussian process surrogate: it is a standard surrogate-based optimization. Also note that the `sbo` option can support optimization over discrete variables (the discrete variables are relaxed) while `ego` cannot.

This is not the same as surrogate_based_global.

gaussian_process

- Keywords Area
- method
- global_evidence
- sbo
- gaussian_process

Use a Gaussian Process surrogate

Specification

Alias: kriging

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	surfpack	Use the Surfpack version of Gaussian Process surrogates

			dakota	Select the built in Gaussian Process surrogate

Description

Dakota uses the Gaussian process (GP) surrogate from Surfpack, which is specified using the surfpack keyword.

A second version of GP surrogates was available in prior versions of Dakota.

For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.

surfpack

- Keywords Area
- method
- global_evidence
- sbo
- gaussian_process
- surfpack

Use the Surfpack version of Gaussian Process surrogates

Specification

Alias: none

Argument(s): none

Description

This keyword specifies the use of the Gaussian process that is incorporated in our surface fitting library called Surfpack.

Several user options are available:

1. Optimization methods:

   Maximum Likelihood Estimation (MLE) is used to find the optimal values of the hyper-parameters governing the trend and correlation functions. By default the global optimization method DIRECT is used for MLE, but other options for the optimization method are available. See optimization_method.

   The total number of evaluations of the likelihood function can be controlled using the `max_trials` keyword followed by a positive integer. Note that the likelihood function does not require running the "truth" model, and is relatively inexpensive to compute.

2. Trend Function:

   The GP models incorporate a parametric trend function whose purpose is to capture large-scale variations. See trend.

3. Correlation Lengths:

   Correlation lengths are usually optimized by Surfpack, however, the user can specify the lengths manually. See correlation_lengths.

4. Ill-conditioning

   One of the major problems in determining the governing values for a Gaussian process or Kriging model is the fact that the correlation matrix can easily become ill-conditioned when there are too many input points close together. Since the predictions from the Gaussian process model involve inverting the correlation matrix, ill-conditioning can lead to poor predictive capability and should be avoided.

   Note that a sufficiently bad sample design could require correlation lengths to be so short that any interpolatory GP model would become inept at extrapolation and interpolation.

   The `surfpack` model handles ill-conditioning internally by default, but behavior can be modified using

5. Gradient Enhanced Kriging (GEK).

   The `use_derivatives` keyword will cause the Surfpack GP to be constructed from a combination of function value and gradient information (if available).

   See notes in the Theory section.

Theory

Gradient Enhanced Kriging

Incorporating gradient information will only be beneficial if accurate and inexpensive derivative information is available, and the derivatives are not infinite or nearly so. Here "inexpensive" means that the cost of evaluating a function value plus gradient is comparable to the cost of evaluating only the function value, for example gradients computed by analytical, automatic differentiation, or continuous adjoint techniques. It is not cost effective to use derivatives computed by finite differences. In tests, GEK models built from finite difference derivatives were also significantly less accurate than those built from analytical derivatives. Note that GEK's correlation matrix tends to have a significantly worse condition number than Kriging for the same sample design.

This issue was addressed by using a pivoted Cholesky factorization of Kriging's correlation matrix (which is a small sub-matrix within GEK's correlation matrix) to rank points by how much unique information they contain. This reordering is then applied to whole points (the function value at a point immediately followed by gradient information at the same point) in GEK's correlation matrix. A standard non-pivoted Cholesky is then applied to the reordered GEK correlation matrix and a bisection search is used to find the last equation that meets the constraint on the (estimate of) condition number. The cost of performing pivoted Cholesky on Kriging's correlation matrix is usually negligible compared to the cost of the non-pivoted Cholesky factorization of GEK's correlation matrix. In tests, it also resulted in more accurate GEK models than when pivoted Cholesky or whole-point-block pivoted Cholesky was performed on GEK's correlation matrix.

?? This correlation function results in a response surface that is C^{∞}-continuous.

dakota

- Keywords Area
- method
- global_evidence
- sbo
- gaussian_process
- dakota

Select the built in Gaussian Process surrogate

Specification

Alias: none
Argument(s): none

Description

A second version of GP surrogates was available in prior versions of Dakota. **For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.**

Historically these models were drastically different, but in Dakota 5.1, they became quite similar. They now differ in that the Surfpack GP has a richer set of features/options and tends to be more accurate than the Dakota version. Due to how the Surfpack GP handles ill-conditioned correlation matrices (which significantly contributes to its greater accuracy), the `Surfpack` GP can be a factor of two or three slower than Dakota's. As of Dakota 5.2, the Surfpack implementation is the default in all contexts except Bayesian calibration.

More details on the `gaussian_process` dakota model can be found in[59].

Dakota's GP deals with ill-conditioning in two ways. First, when it encounters a non-invertible correlation matrix it iteratively increases the size of a "nugget," but in such cases the resulting approximation smooths rather than interpolates the data. Second, it has a `point_selection` option (default off) that uses a greedy algorithm to select a well-spaced subset of points prior to the construction of the GP. In this case, the GP will only interpolate the selected subset. Typically, one should not need point selection in trust-region methods because a small number of points are used to develop a surrogate within each trust region. Point selection is most beneficial when constructing with a large number of points, typically more than order one hundred, though this depends on the number of variables and spacing of the sample points.

This differs from the `point_selection` option of the Dakota GP which initially chooses a well-spaced subset of points and finds the correlation parameters that are most likely for that one subset.

use_derivatives

- Keywords Area
- method
- global_evidence
- sbo
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none
Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

import_points_file

- Keywords Area
- method
- global_evidence
- sbo
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- global_evidence
- sbo
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- global_evidence
- sbo
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- global_evidence
- sbo
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- global_evidence
- sbo
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- global_evidence
- sbo
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

ego

- Keywords Area
- method
- global_evidence
- ego

Use the Efficient Global Optimization method

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		gaussian_process	Use a Gaussian Process surrogate
	Optional		use_derivatives	Use derivative data to construct surrogate models

	Optional	import_points_file	Specify file containing points (variable and response values)
	Optional	export_points_file	Specify file to write out the responses from the surrogate

Description

In the case of `ego`, the efficient global optimization (EGO) method is used to calculate bounds. By default, the Surfpack GP (Kriging) model is used, but the Dakota implementation may be selected instead. If `use_derivatives` is specified the GP model will be built using available derivative data (Surfpack GP only).

See efficient_global for more information.

gaussian_process

- Keywords Area
- method
- global_evidence
- ego
- gaussian_process

Use a Gaussian Process surrogate

Specification

Alias: kriging

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	surfpack	Use the Surfpack version of Gaussian Process surrogates
			dakota	Select the built in Gaussian Process surrogate

Description

Dakota uses the Gaussian process (GP) surrogate from Surfpack, which is specified using the surfpack keyword.

A second version of GP surrogates was available in prior versions of Dakota.

For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.

surfpack

- Keywords Area
- method
- global_evidence
- ego
- gaussian_process
- surfpack

Use the Surfpack version of Gaussian Process surrogates

Specification

Alias: none
Argument(s): none

Description

This keyword specifies the use of the Gaussian process that is incorporated in our surface fitting library called Surfpack.

Several user options are available:

1. Optimization methods:

   Maximum Likelihood Estimation (MLE) is used to find the optimal values of the hyper-parameters governing the trend and correlation functions. By default the global optimization method DIRECT is used for MLE, but other options for the optimization method are available. See optimization_method.

   The total number of evaluations of the likelihood function can be controlled using the `max_trials` keyword followed by a positive integer. Note that the likelihood function does not require running the "truth" model, and is relatively inexpensive to compute.

2. Trend Function:

   The GP models incorporate a parametric trend function whose purpose is to capture large-scale variations. See trend.

3. Correlation Lengths:

   Correlation lengths are usually optimized by Surfpack, however, the user can specify the lengths manually. See correlation_lengths.

4. Ill-conditioning

   One of the major problems in determining the governing values for a Gaussian process or Kriging model is the fact that the correlation matrix can easily become ill-conditioned when there are too many input points close together. Since the predictions from the Gaussian process model involve inverting the correlation matrix, ill-conditioning can lead to poor predictive capability and should be avoided.

   Note that a sufficiently bad sample design could require correlation lengths to be so short that any interpolatory GP model would become inept at extrapolation and interpolation.

   The `surfpack` model handles ill-conditioning internally by default, but behavior can be modified using

5. Gradient Enhanced Kriging (GEK).

   The `use_derivatives` keyword will cause the Surfpack GP to be constructed from a combination of function value and gradient information (if available).

   See notes in the Theory section.

Theory

Gradient Enhanced Kriging

Incorporating gradient information will only be beneficial if accurate and inexpensive derivative information is available, and the derivatives are not infinite or nearly so. Here "inexpensive" means that the cost of evaluating a function value plus gradient is comparable to the cost of evaluating only the function value, for example gradients computed by analytical, automatic differentiation, or continuous adjoint techniques. It is not cost effective to use derivatives computed by finite differences. In tests, GEK models built from finite difference derivatives were also significantly less accurate than those built from analytical derivatives. Note that GEK's correlation matrix tends to have a significantly worse condition number than Kriging for the same sample design.

This issue was addressed by using a pivoted Cholesky factorization of Kriging's correlation matrix (which is a small sub-matrix within GEK's correlation matrix) to rank points by how much unique information they contain. This reordering is then applied to whole points (the function value at a point immediately followed by gradient information at the same point) in GEK's correlation matrix. A standard non-pivoted Cholesky is then applied to the reordered GEK correlation matrix and a bisection search is used to find the last equation that meets the constraint on the (estimate of) condition number. The cost of performing pivoted Cholesky on Kriging's correlation matrix is usually negligible compared to the cost of the non-pivoted Cholesky factorization of GEK's correlation matrix. In tests, it also resulted in more accurate GEK models than when pivoted Cholesky or whole-point-block pivoted Cholesky was performed on GEK's correlation matrix.

?? This correlation function results in a response surface that is C^{∞}-continuous.

dakota

- Keywords Area
- method
- global_evidence
- ego
- gaussian_process
- dakota

Select the built in Gaussian Process surrogate

Specification

Alias: none

Argument(s): none

Description

A second version of GP surrogates was available in prior versions of Dakota. **For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.**

Historically these models were drastically different, but in Dakota 5.1, they became quite similar. They now differ in that the Surfpack GP has a richer set of features/options and tends to be more accurate than the Dakota version. Due to how the Surfpack GP handles ill-conditioned correlation matrices (which significantly contributes to its greater accuracy), the `Surfpack` GP can be a factor of two or three slower than Dakota's. As of Dakota 5.2, the Surfpack implementation is the default in all contexts except Bayesian calibration.

More details on the `gaussian_process` dakota model can be found in[59].

Dakota's GP deals with ill-conditioning in two ways. First, when it encounters a non-invertible correlation matrix it iteratively increases the size of a "nugget," but in such cases the resulting approximation smooths rather than interpolates the data. Second, it has a `point_selection` option (default off) that uses a greedy algorithm to select a well-spaced subset of points prior to the construction of the GP. In this case, the GP will only interpolate the selected subset. Typically, one should not need point selection in trust-region methods because a small number of points are used to develop a surrogate within each trust region. Point selection is most beneficial when constructing with a large number of points, typically more than order one hundred, though this depends on the number of variables and spacing of the sample points.

This differs from the `point_selection` option of the Dakota GP which initially chooses a well-spaced subset of points and finds the correlation parameters that are most likely for that one subset.

use_derivatives

- Keywords Area
- method
- global_evidence
- ego
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none

Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

import_points_file

- Keywords Area
- method

- global_evidence
- ego
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	Description of Group **Group 1**	Dakota Keyword annotated	Dakota Keyword Description (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- global_evidence
- ego
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- global_evidence
- ego
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- global_evidence
- ego
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- global_evidence
- ego
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- global_evidence
- ego
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

ea

- Keywords Area
- method
- global_evidence
- ea

Use an evolutionary algorithm

Specification

Alias: none
Argument(s): none

Description

In this approach, the evolutionary algorithm from Coliny, coliny_ea, is used to perform the interval optimization with no surrogate model involved. Again, this option of `ea` can support interval optimization over discrete variables.

lhs

- Keywords Area
- method
- global_evidence
- lhs

Use LHS method

Specification

Alias: none
Argument(s): none

Description

Use Latin Hypercube Sampling

response_levels

- Keywords Area
- method
- global_evidence
- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_levels` correspond to which response

	Optional	compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- global_evidence
- response_levels

- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- global_evidence
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	probabilities	Compute probabilities
			gen_reliabilities	Compute generalized reliabilities
	Optional		system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- global_evidence
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- global_evidence
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- global_evidence
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional Required*(Choose One)*	Description of Group	Dakota Keyword	Dakota Keyword Description
		Group 1	series	?????
			parallel	?????

series

- Keywords Area
- method
- global_evidence
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

parallel

- Keywords Area
- method
- global_evidence
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

distribution

- Keywords Area
- method
- global_evidence
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	cumulative	Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.

This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- global_evidence
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none

Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- global_evidence
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- global_evidence
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- global_evidence
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- global_evidence
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- global_evidence
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none

Argument(s): INTEGERLIST

Description

See parent page

rng

- Keywords Area
- method
- global_evidence
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional Required*(Choose One)*	Description of Group Group 1	Dakota Keyword mt19937	Dakota Keyword Description Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- global_evidence
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: `http://en.wikipedia.org/wiki/Mersenne_twister`

rnum2

- Keywords Area
- method
- global_evidence
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

samples

- Keywords Area
- method
- global_evidence
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- global_evidence
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- global_evidence
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.60 global_interval_est

- Keywords Area
- method
- global_interval_est

Interval analysis using global optimization methods

Topics

This keyword is related to the topics:

- uncertainty_quantification
- epistemic_uncertainty_quantification_methods
- interval_estimation

Specification

Alias: nond_global_interval_est

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	sbo	Use the surrogate based optimization method
			ego	Use the Efficient Global Optimization method
			ea	Use an evolutionary algorithm
			lhs	Use LHS method
	Optional		rng	Choose a random number generator
	Optional		samples	Specify the number of samples
	Optional		seed	Specify the seed of the random number generator
	Optional		model_pointer	Specify which model block will be used by a method

Description

In the global approach to interval estimation, one uses either a global optimization method or a sampling method to assess the bounds of the responses.

`global_interval_est` allows the user to specify several approaches to calculate interval bounds on the output responses.

- `lhs` - note: this takes the minimum and maximum of the samples as the bounds
- `ego`
- `sbo`
- `ea`

See Also

These keywords may also be of interest:

- global_evidence
- local_evidence
- local_interval_est

sbo

- Keywords Area
- method
- global_interval_est
- sbo

Use the surrogate based optimization method

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		gaussian_process	Use a Gaussian Process surrogate
	Optional		use_derivatives	Use derivative data to construct surrogate models

	Optional	import_points_file	Specify file containing points (variable and response values)
	Optional	export_points_file	Specify file to write out the responses from the surrogate

Description

A surrogate-based optimization method will be used. The surrogate employed in `sbo` is a Gaussian process surrogate.

The main difference between `ego` and the `sbo` approach is the objective function being optimized. `ego` relies on an expected improvement function, while in `sbo`, the optimization proceeds using an evolutionary algorithm (coliny_ea) on the Gaussian process surrogate: it is a standard surrogate-based optimization. Also note that the `sbo` option can support optimization over discrete variables (the discrete variables are relaxed) while `ego` cannot.

This is not the same as surrogate_based_global.

gaussian_process

- Keywords Area
- method
- global_interval_est
- sbo
- gaussian_process

Use a Gaussian Process surrogate

Specification

Alias: kriging

Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** surfpack	**Dakota Keyword Description** Use the Surfpack version of Gaussian Process surrogates
			dakota	Select the built in Gaussian Process surrogate

Description

Dakota uses the Gaussian process (GP) surrogate from Surfpack, which is specified using the surfpack keyword.

A second version of GP surrogates was available in prior versions of Dakota.

For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.

surfpack

- Keywords Area
- method
- global_interval_est
- sbo
- gaussian_process
- surfpack

Use the Surfpack version of Gaussian Process surrogates

Specification

Alias: none
Argument(s): none

Description

This keyword specifies the use of the Gaussian process that is incorporated in our surface fitting library called Surfpack.

Several user options are available:

1. Optimization methods:

   Maximum Likelihood Estimation (MLE) is used to find the optimal values of the hyper-parameters governing the trend and correlation functions. By default the global optimization method DIRECT is used for MLE, but other options for the optimization method are available. See optimization_method.

   The total number of evaluations of the likelihood function can be controlled using the `max_trials` keyword followed by a positive integer. Note that the likelihood function does not require running the "truth" model, and is relatively inexpensive to compute.

2. Trend Function:

   The GP models incorporate a parametric trend function whose purpose is to capture large-scale variations. See trend.

3. Correlation Lengths:

   Correlation lengths are usually optimized by Surfpack, however, the user can specify the lengths manually. See correlation_lengths.

4. Ill-conditioning

   One of the major problems in determining the governing values for a Gaussian process or Kriging model is the fact that the correlation matrix can easily become ill-conditioned when there are too many input points close together. Since the predictions from the Gaussian process model involve inverting the correlation matrix, ill-conditioning can lead to poor predictive capability and should be avoided.

   Note that a sufficiently bad sample design could require correlation lengths to be so short that any interpolatory GP model would become inept at extrapolation and interpolation.

   The `surfpack` model handles ill-conditioning internally by default, but behavior can be modified using

5. Gradient Enhanced Kriging (GEK).

   The `use_derivatives` keyword will cause the Surfpack GP to be constructed from a combination of function value and gradient information (if available).

   See notes in the Theory section.

Theory

Gradient Enhanced Kriging

Incorporating gradient information will only be beneficial if accurate and inexpensive derivative information is available, and the derivatives are not infinite or nearly so. Here "inexpensive" means that the cost of evaluating a function value plus gradient is comparable to the cost of evaluating only the function value, for example gradients computed by analytical, automatic differentiation, or continuous adjoint techniques. It is not cost effective to use derivatives computed by finite differences. In tests, GEK models built from finite difference derivatives were also significantly less accurate than those built from analytical derivatives. Note that GEK's correlation matrix tends to have a significantly worse condition number than Kriging for the same sample design.

This issue was addressed by using a pivoted Cholesky factorization of Kriging's correlation matrix (which is a small sub-matrix within GEK's correlation matrix) to rank points by how much unique information they contain. This reordering is then applied to whole points (the function value at a point immediately followed by gradient information at the same point) in GEK's correlation matrix. A standard non-pivoted Cholesky is then applied to the reordered GEK correlation matrix and a bisection search is used to find the last equation that meets the constraint on the (estimate of) condition number. The cost of performing pivoted Cholesky on Kriging's correlation matrix is usually negligible compared to the cost of the non-pivoted Cholesky factorization of GEK's correlation matrix. In tests, it also resulted in more accurate GEK models than when pivoted Cholesky or whole-point-block pivoted Cholesky was performed on GEK's correlation matrix.

?? This correlation function results in a response surface that is C^{∞}-continuous.

dakota

- Keywords Area
- method
- global_interval_est
- sbo
- gaussian_process
- dakota

Select the built in Gaussian Process surrogate

Specification

Alias: none

Argument(s): none

Description

A second version of GP surrogates was available in prior versions of Dakota. **For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.**

Historically these models were drastically different, but in Dakota 5.1, they became quite similar. They now differ in that the Surfpack GP has a richer set of features/options and tends to be more accurate than the Dakota version. Due to how the Surfpack GP handles ill-conditioned correlation matrices (which significantly contributes to its greater accuracy), the `Surfpack` GP can be a factor of two or three slower than Dakota's. As of Dakota 5.2, the Surfpack implementation is the default in all contexts except Bayesian calibration.

More details on the `gaussian_process` dakota model can be found in[59].

Dakota's GP deals with ill-conditioning in two ways. First, when it encounters a non-invertible correlation matrix it iteratively increases the size of a "nugget," but in such cases the resulting approximation smooths rather than interpolates the data. Second, it has a `point_selection` option (default off) that uses a greedy algorithm to select a well-spaced subset of points prior to the construction of the GP. In this case, the GP will only interpolate the selected subset. Typically, one should not need point selection in trust-region methods because a small number of points are used to develop a surrogate within each trust region. Point selection is most beneficial when constructing with a large number of points, typically more than order one hundred, though this depends on the number of variables and spacing of the sample points.

This differs from the `point_selection` option of the Dakota GP which initially chooses a well-spaced subset of points and finds the correlation parameters that are most likely for that one subset.

use_derivatives

- Keywords Area
- method
- global_interval_est
- sbo
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none

Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

import_points_file

- Keywords Area
- method

- global_interval_est
- sbo
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	Description of Group **Group 1**	Dakota Keyword annotated	Dakota Keyword Description (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- global_interval_est
- sbo
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- global_interval_est
- sbo
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- global_interval_est
- sbo
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional**(*Choose One*)	Description of Group **Group 1**	Dakota Keyword annotated	Dakota Keyword Description (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- global_interval_est
- sbo
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- global_interval_est
- sbo
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

ego

- Keywords Area
- method
- global_interval_est
- ego

Use the Efficient Global Optimization method

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		gaussian_process	Use a Gaussian Process surrogate
	Optional		use_derivatives	Use derivative data to construct surrogate models

	Optional	import_points_file	Specify file containing points (variable and response values)
	Optional	export_points_file	Specify file to write out the responses from the surrogate

Description

In the case of `ego`, the efficient global optimization (EGO) method is used to calculate bounds. By default, the Surfpack GP (Kriging) model is used, but the Dakota implementation may be selected instead. If `use_derivatives` is specified the GP model will be built using available derivative data (Surfpack GP only).

See efficient_global for more information.

gaussian_process

- Keywords Area
- method
- global_interval_est
- ego
- gaussian_process

Use a Gaussian Process surrogate

Specification

Alias: kriging

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	surfpack	Use the Surfpack version of Gaussian Process surrogates
			dakota	Select the built in Gaussian Process surrogate

Description

Dakota uses the Gaussian process (GP) surrogate from Surfpack, which is specified using the surfpack keyword.

A second version of GP surrogates was available in prior versions of Dakota.

For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.

surfpack

- Keywords Area
- method
- global_interval_est
- ego
- gaussian_process
- surfpack

Use the Surfpack version of Gaussian Process surrogates

Specification

Alias: none
Argument(s): none

Description

This keyword specifies the use of the Gaussian process that is incorporated in our surface fitting library called Surfpack.

Several user options are available:

1. Optimization methods:

   Maximum Likelihood Estimation (MLE) is used to find the optimal values of the hyper-parameters governing the trend and correlation functions. By default the global optimization method DIRECT is used for MLE, but other options for the optimization method are available. See optimization_method.

   The total number of evaluations of the likelihood function can be controlled using the `max_trials` keyword followed by a positive integer. Note that the likelihood function does not require running the "truth" model, and is relatively inexpensive to compute.

2. Trend Function:

   The GP models incorporate a parametric trend function whose purpose is to capture large-scale variations. See trend.

3. Correlation Lengths:

   Correlation lengths are usually optimized by Surfpack, however, the user can specify the lengths manually. See correlation_lengths.

4. Ill-conditioning

   One of the major problems in determining the governing values for a Gaussian process or Kriging model is the fact that the correlation matrix can easily become ill-conditioned when there are too many input points close together. Since the predictions from the Gaussian process model involve inverting the correlation matrix, ill-conditioning can lead to poor predictive capability and should be avoided.

   Note that a sufficiently bad sample design could require correlation lengths to be so short that any interpolatory GP model would become inept at extrapolation and interpolation.

   The `surfpack` model handles ill-conditioning internally by default, but behavior can be modified using

5. Gradient Enhanced Kriging (GEK).

   The `use_derivatives` keyword will cause the Surfpack GP to be constructed from a combination of function value and gradient information (if available).

   See notes in the Theory section.

Theory

Gradient Enhanced Kriging

Incorporating gradient information will only be beneficial if accurate and inexpensive derivative information is available, and the derivatives are not infinite or nearly so. Here "inexpensive" means that the cost of evaluating a function value plus gradient is comparable to the cost of evaluating only the function value, for example gradients computed by analytical, automatic differentiation, or continuous adjoint techniques. It is not cost effective to use derivatives computed by finite differences. In tests, GEK models built from finite difference derivatives were also significantly less accurate than those built from analytical derivatives. Note that GEK's correlation matrix tends to have a significantly worse condition number than Kriging for the same sample design.

This issue was addressed by using a pivoted Cholesky factorization of Kriging's correlation matrix (which is a small sub-matrix within GEK's correlation matrix) to rank points by how much unique information they contain. This reordering is then applied to whole points (the function value at a point immediately followed by gradient information at the same point) in GEK's correlation matrix. A standard non-pivoted Cholesky is then applied to the reordered GEK correlation matrix and a bisection search is used to find the last equation that meets the constraint on the (estimate of) condition number. The cost of performing pivoted Cholesky on Kriging's correlation matrix is usually negligible compared to the cost of the non-pivoted Cholesky factorization of GEK's correlation matrix. In tests, it also resulted in more accurate GEK models than when pivoted Cholesky or whole-point-block pivoted Cholesky was performed on GEK's correlation matrix.

?? This correlation function results in a response surface that is C^{∞}-continuous.

dakota

- Keywords Area
- method
- global_interval_est
- ego
- gaussian_process
- dakota

Select the built in Gaussian Process surrogate

Specification

Alias: none

Argument(s): none

Description

A second version of GP surrogates was available in prior versions of Dakota. **For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.**

Historically these models were drastically different, but in Dakota 5.1, they became quite similar. They now differ in that the Surfpack GP has a richer set of features/options and tends to be more accurate than the Dakota version. Due to how the Surfpack GP handles ill-conditioned correlation matrices (which significantly contributes to its greater accuracy), the `Surfpack` GP can be a factor of two or three slower than Dakota's. As of Dakota 5.2, the Surfpack implementation is the default in all contexts except Bayesian calibration.

More details on the `gaussian_process` dakota model can be found in[59].

Dakota's GP deals with ill-conditioning in two ways. First, when it encounters a non-invertible correlation matrix it iteratively increases the size of a "nugget," but in such cases the resulting approximation smooths rather than interpolates the data. Second, it has a `point_selection` option (default off) that uses a greedy algorithm to select a well-spaced subset of points prior to the construction of the GP. In this case, the GP will only interpolate the selected subset. Typically, one should not need point selection in trust-region methods because a small number of points are used to develop a surrogate within each trust region. Point selection is most beneficial when constructing with a large number of points, typically more than order one hundred, though this depends on the number of variables and spacing of the sample points.

This differs from the `point_selection` option of the Dakota GP which initially chooses a well-spaced subset of points and finds the correlation parameters that are most likely for that one subset.

use_derivatives

- Keywords Area
- method
- global_interval_est
- ego
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none

Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

import_points_file

- Keywords Area
- method

- global_interval_est
- ego
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	Description of Group **Group 1**	Dakota Keyword annotated	Dakota Keyword Description (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- global_interval_est
- ego
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- global_interval_est
- ego
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- global_interval_est
- ego
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	Description of Group **Group 1**	Dakota Keyword annotated	Dakota Keyword Description (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- global_interval_est
- ego
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- global_interval_est
- ego
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

ea

- Keywords Area
- method
- global_interval_est
- ea

Use an evolutionary algorithm

Specification

Alias: none
Argument(s): none

Description

In this approach, the evolutionary algorithm from Coliny, coliny_ea, is used to perform the interval optimization with no surrogate model involved. Again, this option of `ea` can support interval optimization over discrete variables.

lhs

- Keywords Area
- method
- global_interval_est
- lhs

Use LHS method

Specification

Alias: none
Argument(s): none

Description

Use Latin Hypercube Sampling

rng

- Keywords Area
- method
- global_interval_est
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	mt19937	Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- global_interval_est
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: http://en.wikipedia.org/wiki/Mersenne_twister

rnum2

- Keywords Area
- method
- global_interval_est
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

samples

- Keywords Area
- method
- global_interval_est
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- global_interval_est
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- global_interval_est
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.61 bayes_calibration

- Keywords Area
- method
- bayes_calibration

Bayesian calibration

Topics

This keyword is related to the topics:

- bayesian_calibration
- package_queso

Specification

Alias: nond_bayes_calibration
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required(*Choose One*)	**Group 1**	queso	Markov Chain Monte Carlo algorithms from the QUESO package

			gpmsa	Gaussian Process Models for Simulation Analysis (GPMSA) Markov Chain Monte Carlo algorithm with Gaussian Process Surrogate
			dream	DREAM (DiffeRential Evolution Adaptive Metropolis)
	Optional		use_derivatives	Use derivative data to construct surrogate models
	Optional		likelihood_scale	Scale the **log-likelihood** function
	Optional		calibrate_sigma	Calibrate the `sigma` term(s) in the likelihood function
	Optional		samples	Specify the number of samples
	Optional		seed	Specify the seed of the random number generator
	Optional		model_pointer	Specify which model block will be used by a method

Description

Currently, we are in the process of incorporating Bayesian calibration methods in Dakota. These methods take prior information on parameter values (in the form of prior distributions) and observational data (e.g. from experiments) and produce posterior distributions on the parameter values. When the computational simulation is then executed with samples from the posterior parameter distributions, the results that are produced are consistent with ("agree with") the experimental data. In the case of calibrating parameters from a computational simulation model, we require a "likelihood function" that specifies the likelihood of observing a particular observation given the model and its associated parameterization. We assume a Gaussian likelihood function currently. The algorithms that produce the posterior distributions on model parameters are Monte Carlo Markov Chain (MCMC) sampling algorithms. MCMC methods require many samples, often tens of thousands, so in the case of model calibration, often emulators of the computational simulation are used. For more details on the algorithms underlying the methods, see the Dakota User's manual.

We have three Bayesian calibration methods under development in Dakota: one called QUESO, one called D-REAM, and one called GPMSA. They are specified with the `bayes_calibration queso`, `bayes_calibration dream`, or `bayes_calibration gpmsa`, respectively. The QUESO method uses components from the QUESO library (Quantification of Uncertainty for Estimation, Simulation, and Optimization) developed at The Uni-

versity of Texas at Austin. We are using a DRAM (Delayed Rejected Adaptive Metropolis) algorithm for the MCMC sampling from the QUESO library. DREAM (DiffeRential Evolution Adaptive Metropolis) is a method that runs multiple different chains simultaneously for global exploration, and automatically tunes the proposal covariance during the process by a self-adaptive randomized subspace sampling. [[87] "Vrugt et al. 2009"]. GPMSA (Gaussian Process Models for Simulation Analysis) is a code that has been developed at Los Alamos National Laboratory. It uses Gaussian process models as part of constructing an emulator for the expensive computational simulation. GPMSA also has extensive features for calibration, such as the capability to include a "model discrepancy" term and the capability to model functional data such as time series data.

As mentioned above, the Bayesian capability in Dakota currently relies on the QUESO library developed by The University of Texas at Austin. This integrated capability is still in prototype form and available to close collaborators of the Dakota team. If you are interested in this capability, contact the Dakota developers at `dakota-developers@development.sandia.gov`.

queso

- Keywords Area
- method
- bayes_calibration
- queso

Markov Chain Monte Carlo algorithms from the QUESO package

Topics

This keyword is related to the topics:

- bayesian_calibration
- package_queso

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		emulator	Use an emulator or surrogate model to evaluate the likelihood function
	Optional		mcmc_type	Type of MCMC algorithm used (dram or multilevel).

	Optional	rng	Choose a random number generator
	Optional	proposal_-covariance_scale	Provide values for the diagonals of the proposal covariance matrix.

Description

For the QUESO method, one can use an emulator in the MCMC sampling. This will greatly improve the speed, since the Monte Carlo Markov Chain will generate thousands of samples on the emulator instead of the real simulation code. However, in the case of fast running evaluations, we recommend the use of no emulator. An emulator may be specified with the keyword `emulator`, followed by a `gaussian_process` emulator, a `pce` emulator (polynomial chaos expansion), or a `sc` emulator (stochastic collocation). For the `gaussian_process` emulator, the user must specify whether to use the `surfpack` or `dakota` version of the Gaussian process. The user can define the number of samples `emulator_samples` from which the emulator should be built. It is also possible to build the Gaussian process from points read in from the `import_points_file` and to export approximation-based sample evaluations using `export_points_file`. For `pce` or `sc`, the user can define a `sparse_grid_level`.

In terms of the MCMC sampling, one can specify the following for the QUESO method: With the `metropolis` type, we have the options `hastings` for a standard Metropolis-Hastings algorithm, or `adaptive` for the adaptive Metropolis in which the covariance of the proposal density is updated adaptively. For the delayed rejection part of the DRAM algorithm, one specifies `rejection`, followed by `standard` (no delayed rejection) or `delayed`. Finally, the user has two scale factors which help control the scaling involved in the problem. The `likelihood_scale` is a number which scales the likelihood by dividing the log of the likelihood (e.g. dividing the sum of squared differences between the experimental data and simulation data or SSE). This is useful for situations with very small likelihoods (e.g. the model is either very far away from the data or there is a lot of data so the likelihood function involves multiplying many likelihoods together, where the SSE term is large and the likelihood becomes very small). In some respects, the `likelihood_scale` can be seen as a normalizing factor for the SSE. If the SSE is large, the `likelihood_scale` should be large. The second factor is a `proposal_covariance_scale`. This allows the user to specify the diagonal values of the proposal covariance matrix. This may be useful when the input variables being calibrated are of different magnitudes: one may want to take a larger step in a direction with a larger magnitude, for example. The `proposal_covariance_scale` should be followed by a list of the size of the number of parameters being calibrated, with values for the proposal covariance diagonal for each. Finally, we offer the option to calibrate the sigma terms with the `calibrate_sigma` flag. The sigma terms refer to the error associated with the Gaussian process: sigma is used in the likelihood calculation. If experimental measurement error is available to inform sigma, that is very useful, but often measurement uncertainty is not available. Note that if `calibrate_sigma` is specified, a separate sigma term will be calibrated for each calibration term. Thus, if there are 50 calibration terms (e.g. experimental points against which we are trying to match the model), 50 sigma values will be added to the calibration process. Calibration of the sigma values is turned off by default: only the design parameters are calibrated in default mode.

emulator

- Keywords Area
- method
- bayes_calibration
- queso

- emulator

Use an emulator or surrogate model to evaluate the likelihood function

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	gaussian_process	Use a Gaussian Process surrogate
			pce	Use a Polynomial Chaos Expansion surrogate model
			sc	Use Stochastic Collocation to generate a polynomial surrogate model

Description

This keyword describes the type of emulator used when calculating the likelihood function for the Bayesian calibration. The emulator can be a Gaussian process, polynomial chaos expansion, or stochastic collocation.

gaussian_process

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- gaussian_process

Use a Gaussian Process surrogate

Specification

Alias: kriging
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Required*(Choose One)*	**Group 1**	surfpack	Use the Surfpack version of Gaussian Process surrogates
			dakota	Select the built in Gaussian Process surrogate
	Optional		emulator_samples	Number of data points used to train the surrogate model or emulator
	Optional		import_points_file	Specify file containing points (variable and response values)
	Optional		export_points_file	Specify file to write out the responses from the surrogate

Description

Dakota uses the Gaussian process (GP) surrogate from Surfpack, which is specified using the surfpack keyword.

A second version of GP surrogates was available in prior versions of Dakota.

For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.

surfpack

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- gaussian_process
- surfpack

Use the Surfpack version of Gaussian Process surrogates

Specification

Alias: none

Argument(s): none

Description

This keyword specifies the use of the Gaussian process that is incorporated in our surface fitting library called Surfpack.

Several user options are available:

1. Optimization methods:

   Maximum Likelihood Estimation (MLE) is used to find the optimal values of the hyper-parameters governing the trend and correlation functions. By default the global optimization method DIRECT is used for MLE, but other options for the optimization method are available. See optimization_method.

   The total number of evaluations of the likelihood function can be controlled using the `max_trials` keyword followed by a positive integer. Note that the likelihood function does not require running the "truth" model, and is relatively inexpensive to compute.

2. Trend Function:

   The GP models incorporate a parametric trend function whose purpose is to capture large-scale variations. See trend.

3. Correlation Lengths:

   Correlation lengths are usually optimized by Surfpack, however, the user can specify the lengths manually. See correlation_lengths.

4. Ill-conditioning

   One of the major problems in determining the governing values for a Gaussian process or Kriging model is the fact that the correlation matrix can easily become ill-conditioned when there are too many input points close together. Since the predictions from the Gaussian process model involve inverting the correlation matrix, ill-conditioning can lead to poor predictive capability and should be avoided.

   Note that a sufficiently bad sample design could require correlation lengths to be so short that any interpolatory GP model would become inept at extrapolation and interpolation.

   The `surfpack` model handles ill-conditioning internally by default, but behavior can be modified using

5. Gradient Enhanced Kriging (GEK).

   The `use_derivatives` keyword will cause the Surfpack GP to be constructed from a combination of function value and gradient information (if available).

   See notes in the Theory section.

Theory

Gradient Enhanced Kriging

Incorporating gradient information will only be beneficial if accurate and inexpensive derivative information is available, and the derivatives are not infinite or nearly so. Here "inexpensive" means that the cost of evaluating a function value plus gradient is comparable to the cost of evaluating only the function value, for example gradients computed by analytical, automatic differentiation, or continuous adjoint techniques. It is not cost effective to use derivatives computed by finite differences. In tests, GEK models built from finite difference derivatives were also significantly less accurate than those built from analytical derivatives. Note that GEK's correlation matrix tends to have a significantly worse condition number than Kriging for the same sample design.

This issue was addressed by using a pivoted Cholesky factorization of Kriging's correlation matrix (which is a small sub-matrix within GEK's correlation matrix) to rank points by how much unique information they contain. This reordering is then applied to whole points (the function value at a point immediately followed by gradient

information at the same point) in GEK's correlation matrix. A standard non-pivoted Cholesky is then applied to the reordered GEK correlation matrix and a bisection search is used to find the last equation that meets the constraint on the (estimate of) condition number. The cost of performing pivoted Cholesky on Kriging's correlation matrix is usually negligible compared to the cost of the non-pivoted Cholesky factorization of GEK's correlation matrix. In tests, it also resulted in more accurate GEK models than when pivoted Cholesky or whole-point-block pivoted Cholesky was performed on GEK's correlation matrix.

?? This correlation function results in a response surface that is C^{∞}-continuous.

dakota

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- gaussian_process
- dakota

Select the built in Gaussian Process surrogate

Specification

Alias: none

Argument(s): none

Description

A second version of GP surrogates was available in prior versions of Dakota. **For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.**

Historically these models were drastically different, but in Dakota 5.1, they became quite similar. They now differ in that the Surfpack GP has a richer set of features/options and tends to be more accurate than the Dakota version. Due to how the Surfpack GP handles ill-conditioned correlation matrices (which significantly contributes to its greater accuracy), the `Surfpack` GP can be a factor of two or three slower than Dakota's. As of Dakota 5.2, the Surfpack implementation is the default in all contexts except Bayesian calibration.

More details on the `gaussian_process` dakota model can be found in[59].

Dakota's GP deals with ill-conditioning in two ways. First, when it encounters a non-invertible correlation matrix it iteratively increases the size of a "nugget," but in such cases the resulting approximation smooths rather than interpolates the data. Second, it has a `point_selection` option (default off) that uses a greedy algorithm to select a well-spaced subset of points prior to the construction of the GP. In this case, the GP will only interpolate the selected subset. Typically, one should not need point selection in trust-region methods because a small number of points are used to develop a surrogate within each trust region. Point selection is most beneficial when constructing with a large number of points, typically more than order one hundred, though this depends on the number of variables and spacing of the sample points.

This differs from the `point_selection` option of the Dakota GP which initially chooses a well-spaced subset of points and finds the correlation parameters that are most likely for that one subset.

emulator_samples

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- gaussian_process
- emulator_samples

Number of data points used to train the surrogate model or emulator

Specification

Alias: none
Argument(s): INTEGER

Description

This keyword refers to the number of build points or training points used to construct a Gaussian process emulator. If the user specifies a number of `emulator_samples` that is less than the minimum number of points required to build the GP surrogate, Dakota will augment the samples to obtain the minimum required.

import_points_file

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- gaussian_process
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- gaussian_process
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- gaussian_process
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- gaussian_process
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- gaussian_process
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- gaussian_process
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

pce

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- pce

Use a Polynomial Chaos Expansion surrogate model

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		sparse_grid_level	Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Description

This specifies the use of a polynomial chaos expansion (PCE) model to be used as a surrogate model in the likelihood calculations used in the Bayesian computations. When using PCE as a surrogate within the Bayesian framework, the build points (training points) for the PCE are constructed from a sparse grid.

See Also

These keywords may also be of interest:

- polynomial_chaos

sparse_grid_level

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- pce

- sparse_grid_level

Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Specification

Alias: none
Argument(s): INTEGERLIST

Description

Multi-dimensional integration by the Smolyak sparse grid method (specified with sparse_grid_level and, optionally, dimension_preference). The underlying one-dimensional integration rules are the same as for the tensor-product quadrature case; however, the default rule selection is nested for sparse grids (Genz-Keister for normals/transformed normals and Gauss-Patterson for uniforms/transformed uniforms). This default can be overridden with an explicit non_nested specification (resulting in Gauss-Hermite for normals/transformed normals and Gauss-Legendre for uniforms/transformed uniforms). As for tensor quadrature, the dimension_preference specification enables the use of anisotropic sparse grids (refer to the PCE description in the User's Manual for the anisotropic index set constraint definition). Similar to anisotropic tensor grids, the dimension with greatest preference will have resolution at the full sparse_grid_level and all other dimension resolutions will be reduced in proportion to their reduced preference. For PCE with either isotropic or anisotropic sparse grids, a summation of tensor-product expansions is used, where each anisotropic tensor-product quadrature rule underlying the sparse grid construction results in its own anisotropic tensor-product expansion as described in case 1. These anisotropic tensor-product expansions are summed into a sparse PCE using the standard Smolyak summation (again, refer to the User's Manual for additional details). As for quadrature_order, the sparse_grid_level specification admits an array input for enabling specification of multiple grid resolutions used by certain advanced solution methodologies.

This keyword can be used when using sparse grid integration to calculate PCE coefficients or when generating samples for sparse grid collocation.

sc

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- sc

Use Stochastic Collocation to generate a polynomial surrogate model

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		sparse_grid_level	Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Description

This specifies the use of a stochastic collocation (SC) model to be used as a surrogate model in the likelihood calculations used in the Bayesian computations. When using SC as a surrogate within the Bayesian framework, the build points (training points) for the stochastic collocation are constructed from a sparse grid.

See Also

These keywords may also be of interest:

- stoch_collocation

sparse_grid_level

- Keywords Area
- method
- bayes_calibration
- queso
- emulator
- sc
- sparse_grid_level

Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Specification

Alias: none
Argument(s): INTEGERLIST

Description

Multi-dimensional integration by the Smolyak sparse grid method (specified with sparse_grid_level and, optionally, dimension_preference). The underlying one-dimensional integration rules are the same as for the tensor-product quadrature case; however, the default rule selection is nested for sparse grids (Genz-Keister for normals/transformed normals and Gauss-Patterson for uniforms/transformed uniforms). This default can be overridden with an explicit non_nested specification (resulting in Gauss-Hermite for normals/transformed normals

and Gauss-Legendre for uniforms/transformed uniforms). As for tensor quadrature, the dimension_preference specification enables the use of anisotropic sparse grids (refer to the PCE description in the User's Manual for the anisotropic index set constraint definition). Similar to anisotropic tensor grids, the dimension with greatest preference will have resolution at the full sparse_grid_level and all other dimension resolutions will be reduced in proportion to their reduced preference. For PCE with either isotropic or anisotropic sparse grids, a summation of tensor-product expansions is used, where each anisotropic tensor-product quadrature rule underlying the sparse grid construction results in its own anisotropic tensor-product expansion as described in case 1. These anisotropic tensor-product expansions are summed into a sparse PCE using the standard Smolyak summation (again, refer to the User's Manual for additional details). As for quadrature_order, the sparse_grid_level specification admits an array input for enabling specification of multiple grid resolutions used by certain advanced solution methodologies.

This keyword can be used when using sparse grid integration to calculate PCE coefficients or when generating samples for sparse grid collocation.

mcmc_type

- Keywords Area
- method
- bayes_calibration
- queso
- mcmc_type

Type of MCMC algorithm used (dram or multilevel).

Topics

This keyword is related to the topics:

- bayesian_calibration

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	dram	MCMC type
			multilevel	Type of MCMC algorithm

Description

The type of Markov Chain Monte Carlo used. There are two types currently: DRAM (Delayed Rejection Adaptive Metropolis) and multi-level. DRAM is described in [[39] "Haario et al. 2006"]. The multilevel algorithm is described in [[71] "Prudencio and Cheung, 2012"].

dram

- Keywords Area
- method
- bayes_calibration
- queso
- mcmc_type
- dram

MCMC type

Topics

This keyword is related to the topics:

- bayesian_calibration

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		rejection	Select the rejection type used in the Markov Chain Monte Carlo algorithm
	Optional		metropolis	Choose the sampler for Markov Chain Monte Carlo

Description

Use of the Delayed Rejection Adaptive Metropolis algorithm (dram).

rejection

- Keywords Area
- method
- bayes_calibration
- queso
- mcmc_type
- dram
- rejection

Select the rejection type used in the Markov Chain Monte Carlo algorithm

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	standard	Select standard rejection
			delayed	Select delayed rejection

Description

Specification of the rejection type used in the Markov Chain Monte Carlo algorithm. The rejection type has two options: `standard` or `delayed`. Standard rejection involves no delayed rejection, while delayed involves a delay in rejecting samples in the MCMC chain. Delayed rejection plus adaptive metropolis is the DRAM algorithm.

standard

- Keywords Area
- method
- bayes_calibration
- queso
- mcmc_type
- dram
- rejection
- standard

Select standard rejection

Specification

Alias: none

Argument(s): none

Description

The use of a standard rejection approach in the Monte Carlo Markov Chain algorithm.

delayed

- Keywords Area
- method
- bayes_calibration
- queso
- mcmc_type
- dram
- rejection
- delayed

Select delayed rejection

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The use of a delayed approach in rejecting points in the Markov Chain Monte Carlo algorithm.

metropolis

- Keywords Area
- method
- bayes_calibration
- queso
- mcmc_type
- dram
- metropolis

Choose the sampler for Markov Chain Monte Carlo

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	hastings	Select the Metropolis--Hastings sampler
			adaptive	Type of Metropolis algorithm

Description

The specification of the algorithm which governs the Markov Chain Monte Carlo (MCMC) algorithm used. There currently are two options for metropolis:

- the option `hastings` specifies a standard Metropolis-Hastings algorithm
- the option `adaptive` specifies the adaptive Metropolis in which the covariance of the proposal density is updated adaptively.

hastings

- Keywords Area
- method
- bayes_calibration
- queso
- mcmc_type
- dram
- metropolis
- hastings

Select the Metropolis-Hastings sampler

Specification

Alias: none
Argument(s): none

Description

The use of the Metropolis-Hastings algorithm in the Markov Chain Monte Carlo algorithm.

adaptive

- Keywords Area
- method
- bayes_calibration
- queso
- mcmc_type
- dram
- metropolis
- adaptive

Type of Metropolis algorithm

Topics

This keyword is related to the topics:

- bayesian_calibration

Specification

Alias: none
Argument(s): none

Description

Type of Metropolis algorithm used in DRAM.

multilevel

- Keywords Area
- method
- bayes_calibration
- queso
- mcmc_type
- multilevel

Type of MCMC algorithm

Topics

This keyword is related to the topics:

- bayesian_calibration

Specification

Alias: none
Argument(s): none

Description

Use of the multilevel MCMC approach described in [[71] "Prudencio and Cheung, 2012"].

rng

- Keywords Area
- method
- bayes_calibration
- queso
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** mt19937	**Dakota Keyword Description** Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- bayes_calibration
- queso
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: http://en.wikipedia.org/wiki/Mersenne_twister

rnum2

- Keywords Area
- method
- bayes_calibration
- queso
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

proposal_covariance_scale

- Keywords Area
- method
- bayes_calibration
- queso
- proposal_covariance_scale

Provide values for the diagonals of the proposal covariance matrix.

Topics

This keyword is related to the topics:

- bayesian_calibration
- problem

Specification

Alias: none
Argument(s): REALLIST

Description

This keyword allows the user to specify the proposal covariance matrix for the Monte Carlo Markov Chain algorithm used in Bayesian calibration. The `proposal_covariance_scale` should be followed by a number of terms equal to the number of variables being calibrated. Each particular value represents an estimate of the variance in that particular parameter. This proposal covariance is taken to be the covariance governing a multivariate normal distribution. Samples are generated from this distribution in the MCMC process. Thus, this proposal covariance is critical in determining how big the step sizes are in the MCMC chain.

This keyword scales the values on the diagonal of the proposal covariance and may be useful when the input variables being calibrated are of different magnitudes: one may want to take a larger step in a direction with a larger magnitude, for example. The proposal covariance is a list of reals, one for each variable.

gpmsa

- Keywords Area
- method
- bayes_calibration
- gpmsa

Gaussian Process Models for Simulation Analysis (GPMSA) Markov Chain Monte Carlo algorithm with Gaussian Process Surrogate

Topics

This keyword is related to the topics:

- package_queso
- bayesian_calibration

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		emulator_samples	Number of data points used to train the surrogate model or emulator
	Optional		import_points_file	Specify file containing points (variable and response values)
	Optional		export_points_file	Specify file to write out the responses from the surrogate
	Optional		mcmc_type	
	Optional		rng	Choose a random number generator
	Optional		proposal_-covariance_scale	Provide values for the diagonals of the proposal covariance matrix.

Description

GPMSA (Gaussian Process Models for Simulation Analysis) is another approach that provides the capability for Bayesian calibration. A key part of GPMSA is the construction of an emulator from simulation runs collected at various settings of input parameters. The emulator is a statistical model of the system response, and it is used to incorporate the observational data to improve system predictions and constrain or calibrate the unknown parameters. The GPMSA code draws heavily on the theory developed in the seminal Bayesian calibration paper by Kennedy and O'Hagan [[56] "Kennedy and O'Hagan 2001"]. The particular approach in GPMSA has been developed by the Los Alamos group and document in [[48] "Higdon et al. 2008"]. GPMSA uses Gaussian process models in the emulation, but the emulator is actually a set of basis functions (e.g. from a singular value decomposition) which have GPs as the coefficients.

For the GPMSA method, one can define the number of samples which will be used in construction of the emulator, `emulator_samples`. The emulator involves Gaussian processes in GPMSA, so the user does not specify anything about emulator type. At this point, the only controls active for GPMSA are `emulator_samples`, `seed` and `rng`, and `samples` (the number of MCMC samples). NOTE: the GPMSA method is in a very preliminary, prototype state at this time.

emulator_samples

- Keywords Area
- method
- bayes_calibration
- gpmsa
- emulator_samples

Number of data points used to train the surrogate model or emulator

Specification

Alias: none
Argument(s): INTEGER

Description

This keyword refers to the number of build points or training points used to construct a Gaussian process emulator. If the user specifies a number of `emulator_samples` that is less than the minimum number of points required to build the GP surrogate, Dakota will augment the samples to obtain the minimum required.

import_points_file

- Keywords Area
- method
- bayes_calibration
- gpmsa
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none
Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- bayes_calibration
- gpmsa
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- bayes_calibration
- gpmsa
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- bayes_calibration
- gpmsa
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional(*Choose One*)	**Group 1**	annotated	(Default) Specified file is in the annotated format

			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- bayes_calibration
- gpmsa
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- bayes_calibration
- gpmsa
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

mcmc_type

- Keywords Area
- method
- bayes_calibration
- gpmsa
- mcmc_type

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required(*Choose One*)	**Group 1**	dram multilevel	

dram

- Keywords Area
- method
- bayes_calibration
- gpmsa
- mcmc_type
- dram

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		rejection	Select the rejection type used in the Markov Chain Monte Carlo algorithm

	Optional	metropolis	Choose the sampler for Markov Chain Monte Carlo

rejection

- Keywords Area
- method
- bayes_calibration
- gpmsa
- mcmc_type
- dram
- rejection

Select the rejection type used in the Markov Chain Monte Carlo algorithm

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** standard	**Dakota Keyword Description** Select standard rejection
			delayed	Select delayed rejection

Description

Specification of the rejection type used in the Markov Chain Monte Carlo algorithm. The rejection type has two options: `standard` or `delayed`. Standard rejection involves no delayed rejection, while delayed involves a delay in rejecting samples in the MCMC chain. Delayed rejection plus adaptive metropolis is the DRAM algorithm.

standard

- Keywords Area
- method
- bayes_calibration
- gpmsa
- mcmc_type
- dram

- rejection
- standard

Select standard rejection

Specification

Alias: none
Argument(s): none

Description

The use of a standard rejection approach in the Monte Carlo Markov Chain algorithm.

delayed

- Keywords Area
- method
- bayes_calibration
- gpmsa
- mcmc_type
- dram
- rejection
- delayed

Select delayed rejection

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The use of a delayed approach in rejecting points in the Markov Chain Monte Carlo algorithm.

metropolis

- Keywords Area
- method
- bayes_calibration
- gpmsa
- mcmc_type
- dram
- metropolis

Choose the sampler for Markov Chain Monte Carlo

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	hastings	Select the Metropolis--Hastings sampler
			adaptive	Select the adaptive Metropolis sampler

Description

The specification of the algorithm which governs the Markov Chain Monte Carlo (MCMC) algorithm used. There currently are two options for metropolis:

- the option `hastings` specifies a standard Metropolis-Hastings algorithm
- the option `adaptive` specifies the adaptive Metropolis in which the covariance of the proposal density is updated adaptively.

hastings

- Keywords Area
- method
- bayes_calibration
- gpmsa
- mcmc_type
- dram
- metropolis
- hastings

Select the Metropolis-Hastings sampler

Specification

Alias: none
Argument(s): none

Description

The use of the Metropolis-Hastings algorithm in the Markov Chain Monte Carlo algorithm.

adaptive

- Keywords Area
- method
- bayes_calibration
- gpmsa
- mcmc_type
- dram
- metropolis
- adaptive

Select the adaptive Metropolis sampler

Specification

Alias: none
Argument(s): none

Description

The use of an adaptive Metropolis algorithm in the Markov Chain Monte Carlo algorithm, involving adaptive updates to the covariance matrix.

multilevel

- Keywords Area
- method
- bayes_calibration
- gpmsa
- mcmc_type
- multilevel

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): none

rng

- Keywords Area
- method
- bayes_calibration
- gpmsa
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** mt19937	**Dakota Keyword Description** Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- bayes_calibration
- gpmsa
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: http://en.wikipedia.org/wiki/Mersenne_twister

rnum2

- Keywords Area
- method
- bayes_calibration
- gpmsa
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

proposal_covariance_scale

- Keywords Area
- method
- bayes_calibration
- gpmsa
- proposal_covariance_scale

Provide values for the diagonals of the proposal covariance matrix.

Topics

This keyword is related to the topics:

- bayesian_calibration
- problem

Specification

Alias: none
Argument(s): REALLIST

Description

This keyword allows the user to specify the proposal covariance matrix for the Monte Carlo Markov Chain algorithm used in Bayesian calibration. The `proposal_covariance_scale` should be followed by a number of terms equal to the number of variables being calibrated. Each particular value represents an estimate of the variance in that particular parameter. This proposal covariance is taken to be the covariance governing a multivariate normal distribution. Samples are generated from this distribution in the MCMC process. Thus, this proposal covariance is critical in determining how big the step sizes are in the MCMC chain.

This keyword scales the values on the diagonal of the proposal covariance and may be useful when the input variables being calibrated are of different magnitudes: one may want to take a larger step in a direction with a larger magnitude, for example. The proposal covariance is a list of reals, one for each variable.

dream

- Keywords Area
- method
- bayes_calibration
- dream

DREAM (DiffeRential Evolution Adaptive Metropolis)

Topics

This keyword is related to the topics:

- bayesian_calibration

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Optional	chains	Number of chains in DREAM
	Optional	num_cr	Number of candidate points for each crossover.
	Optional	crossover_chain_-pairs	Number of chains used in crossover.
	Optional	gr_threshold	Convergence tolerance for the Gelman-Rubin statistic
	Optional	jump_step	Number of generations a long jump step is taken
	Optional	emulator	Use an emulator or surrogate model to evaluate the likelihood function

Description

The DiffeRential Evolution Adaptive Metropolis algorithm is described in [[87] "Vrugt et al. 2009"]. For the DREAM method, one can define the number of chains used with `chains`. The total number of generations per chain in DREAM is the number of samples (`samples`) divided by the number of chains (`chains`). The minimum number of chains is three. The number of chains randomly selected to be used in the crossover each time a crossover occurs is `crossover_chain_pairs`. There is an extra adaptation during burn-in, in which DREAM estimates a distribution of crossover probabilities that favors large jumps over smaller ones in each of the chains. Normalization is required to ensure that all of the input dimensions contribute equally. In this process, a discrete number of candidate points for each crossover value is generated. This parameter is `num_cr`. The `gr_threshold` is the convergence tolerance for the Gelman-Rubin statistic which will govern the convergence of the multiple chain process. The integer `jump_step` forces a long jump every `jump_step` generations. For more details about these parameters, see [[87] "Vrugt et al. 2009"].

chains

- Keywords Area
- method
- bayes_calibration
- dream
- chains

Number of chains in DREAM

Topics

This keyword is related to the topics:

- bayesian_calibration

Specification

Alias: none
Argument(s): INTEGER

Description

Number of chains in DREAM

num_cr

- Keywords Area
- method
- bayes_calibration
- dream
- num_cr

Number of candidate points for each crossover.

Topics

This keyword is related to the topics:

- bayesian_calibration

Specification

Alias: none
Argument(s): INTEGER

Description

In DREAM, there is an extra adaptation during burn-in, in which DREAM estimates a distribution of crossover probabilities that favors large jumps over smaller ones in each of the chains. Normalization is required to ensure that all of the input dimensions contribute equally. In this process, a discrete number of candidate points for each crossover value is generated. This parameter is num_cr.

crossover_chain_pairs

- Keywords Area
- method
- bayes_calibration
- dream
- crossover_chain_pairs

Number of chains used in crossover.

Topics

This keyword is related to the topics:

- bayesian_calibration

Specification

Alias: none
Argument(s): INTEGER

Description

The number of chains randomly selected to be used in the crossover each time a crossover occurs.

gr_threshold

- Keywords Area
- method
- bayes_calibration
- dream
- gr_threshold

Convergence tolerance for the Gelman-Rubin statistic

Topics

This keyword is related to the topics:

- bayesian_calibration

Specification

Alias: none
Argument(s): REAL

Description

The gr_threshold is the convergence tolerance for the Gelman-Rubin statistic which will govern the convergence of the multiple chain process.

jump_step

- Keywords Area
- method
- bayes_calibration
- dream
- jump_step

Number of generations a long jump step is taken

Topics

This keyword is related to the topics:

- bayesian_calibration

Specification

Alias: none
Argument(s): INTEGER

Description

The integer jump_step forces a long jump every jump_step generations in DREAM.

emulator

- Keywords Area
- method
- bayes_calibration
- dream
- emulator

Use an emulator or surrogate model to evaluate the likelihood function

Specification

Alias: none
Argument(s): none

	Required/-Optional	**Description of Group**	**Dakota Keyword**	**Dakota Keyword Description**
	Required*(Choose One)*	**Group 1**	gaussian_process	Use a Gaussian Process surrogate
			pce	Use a polynomial chaos expansion as a surrogate for the truth model when performing Bayesian inference.

			sc	Use a sparse grid as a surrogate for the truth model when performing Bayesian inference.

Description

This keyword describes the type of emulator used when calculating the likelihood function for the Bayesian calibration. The emulator can be a Gaussian process, polynomial chaos expansion, or stochastic collocation.

gaussian_process

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- gaussian_process

Use a Gaussian Process surrogate

Specification

Alias: kriging

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	surfpack	Use the Surfpack version of Gaussian Process surrogates
			dakota	Select the built in Gaussian Process surrogate
	Optional		emulator_samples	Number of data points used to train the surrogate model or emulator

	Optional	import_points_file	Specify file containing points (variable and response values)
	Optional	export_points_file	Specify file to write out the responses from the surrogate

Description

Dakota uses the Gaussian process (GP) surrogate from Surfpack, which is specified using the surfpack keyword.

A second version of GP surrogates was available in prior versions of Dakota.

For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.

surfpack

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- gaussian_process
- surfpack

Use the Surfpack version of Gaussian Process surrogates

Specification

Alias: none

Argument(s): none

Description

This keyword specifies the use of the Gaussian process that is incorporated in our surface fitting library called Surfpack.

Several user options are available:

1. Optimization methods:

   Maximum Likelihood Estimation (MLE) is used to find the optimal values of the hyper-parameters governing the trend and correlation functions. By default the global optimization method DIRECT is used for MLE, but other options for the optimization method are available. See optimization_method.

   The total number of evaluations of the likelihood function can be controlled using the `max_trials` keyword followed by a positive integer. Note that the likelihood function does not require running the "truth" model, and is relatively inexpensive to compute.

2. Trend Function:

   The GP models incorporate a parametric trend function whose purpose is to capture large-scale variations. See trend.

3. Correlation Lengths:

   Correlation lengths are usually optimized by Surfpack, however, the user can specify the lengths manually. See correlation_lengths.

4. Ill-conditioning

   One of the major problems in determining the governing values for a Gaussian process or Kriging model is the fact that the correlation matrix can easily become ill-conditioned when there are too many input points close together. Since the predictions from the Gaussian process model involve inverting the correlation matrix, ill-conditioning can lead to poor predictive capability and should be avoided.

   Note that a sufficiently bad sample design could require correlation lengths to be so short that any interpolatory GP model would become inept at extrapolation and interpolation.

   The `surfpack` model handles ill-conditioning internally by default, but behavior can be modified using

5. Gradient Enhanced Kriging (GEK).

   The `use_derivatives` keyword will cause the Surfpack GP to be constructed from a combination of function value and gradient information (if available).

   See notes in the Theory section.

Theory

Gradient Enhanced Kriging

Incorporating gradient information will only be beneficial if accurate and inexpensive derivative information is available, and the derivatives are not infinite or nearly so. Here "inexpensive" means that the cost of evaluating a function value plus gradient is comparable to the cost of evaluating only the function value, for example gradients computed by analytical, automatic differentiation, or continuous adjoint techniques. It is not cost effective to use derivatives computed by finite differences. In tests, GEK models built from finite difference derivatives were also significantly less accurate than those built from analytical derivatives. Note that GEK's correlation matrix tends to have a significantly worse condition number than Kriging for the same sample design.

This issue was addressed by using a pivoted Cholesky factorization of Kriging's correlation matrix (which is a small sub-matrix within GEK's correlation matrix) to rank points by how much unique information they contain. This reordering is then applied to whole points (the function value at a point immediately followed by gradient information at the same point) in GEK's correlation matrix. A standard non-pivoted Cholesky is then applied to the reordered GEK correlation matrix and a bisection search is used to find the last equation that meets the constraint on the (estimate of) condition number. The cost of performing pivoted Cholesky on Kriging's correlation matrix is usually negligible compared to the cost of the non-pivoted Cholesky factorization of GEK's correlation matrix. In tests, it also resulted in more accurate GEK models than when pivoted Cholesky or whole-point-block pivoted Cholesky was performed on GEK's correlation matrix.

?? This correlation function results in a response surface that is C^{∞}-continuous.

dakota

- Keywords Area
- method

- bayes_calibration
- dream
- emulator
- gaussian_process
- dakota

Select the built in Gaussian Process surrogate

Specification

Alias: none
Argument(s): none

Description

A second version of GP surrogates was available in prior versions of Dakota. **For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.**

Historically these models were drastically different, but in Dakota 5.1, they became quite similar. They now differ in that the Surfpack GP has a richer set of features/options and tends to be more accurate than the Dakota version. Due to how the Surfpack GP handles ill-conditioned correlation matrices (which significantly contributes to its greater accuracy), the `Surfpack` GP can be a factor of two or three slower than Dakota's. As of Dakota 5.2, the Surfpack implementation is the default in all contexts except Bayesian calibration.

More details on the `gaussian_process` dakota model can be found in[59].

Dakota's GP deals with ill-conditioning in two ways. First, when it encounters a non-invertible correlation matrix it iteratively increases the size of a "nugget," but in such cases the resulting approximation smooths rather than interpolates the data. Second, it has a `point_selection` option (default off) that uses a greedy algorithm to select a well-spaced subset of points prior to the construction of the GP. In this case, the GP will only interpolate the selected subset. Typically, one should not need point selection in trust-region methods because a small number of points are used to develop a surrogate within each trust region. Point selection is most beneficial when constructing with a large number of points, typically more than order one hundred, though this depends on the number of variables and spacing of the sample points.

This differs from the `point_selection` option of the Dakota GP which initially chooses a well-spaced subset of points and finds the correlation parameters that are most likely for that one subset.

emulator_samples

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- gaussian_process
- emulator_samples

Number of data points used to train the surrogate model or emulator

Specification

Alias: none
Argument(s): INTEGER

Description

This keyword refers to the number of build points or training points used to construct a Gaussian process emulator. If the user specifies a number of `emulator_samples` that is less than the minimum number of points required to build the GP surrogate, Dakota will augment the samples to obtain the minimum required.

import_points_file

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- gaussian_process
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- gaussian_process
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- gaussian_process
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- gaussian_process
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- gaussian_process
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- gaussian_process
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

pce

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- pce

Use a polynomial chaos expansion as a surrogate for the truth model when performing Bayesian inference.

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		sparse_grid_level	Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Description

Use a polynomial chaos expansion as a surrogate for the truth model when performing Bayesian inference.

sparse_grid_level

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- pce
- sparse_grid_level

Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Specification

Alias: none
Argument(s): INTEGERLIST

Description

Multi-dimensional integration by the Smolyak sparse grid method (specified with sparse_grid_level and, optionally, dimension_preference). The underlying one-dimensional integration rules are the same as for the tensor-product quadrature case; however, the default rule selection is nested for sparse grids (Genz-Keister for normals/transformed normals and Gauss-Patterson for uniforms/transformed uniforms). This default can be overridden with an explicit non_nested specification (resulting in Gauss-Hermite for normals/transformed normals and Gauss-Legendre for uniforms/transformed uniforms). As for tensor quadrature, the dimension_preference specification enables the use of anisotropic sparse grids (refer to the PCE description in the User's Manual for the anisotropic index set constraint definition). Similar to anisotropic tensor grids, the dimension with greatest preference will have resolution at the full sparse_grid_level and all other dimension resolutions will be reduced in proportion to their reduced preference. For PCE with either isotropic or anisotropic sparse grids, a summation of tensor-product expansions is used, where each anisotropic tensor-product quadrature rule underlying the sparse grid construction results in its own anisotropic tensor-product expansion as described in case 1. These anisotropic tensor-product expansions are summed into a sparse PCE using the standard Smolyak summation (again, refer to the User's Manual for additional details). As for quadrature_order, the sparse_grid_level specification admits an array input for enabling specification of multiple grid resolutions used by certain advanced solution methodologies.

This keyword can be used when using sparse grid integration to calculate PCE coefficients or when generating samples for sparse grid collocation.

sc

- Keywords Area
- method
- bayes_calibration
- dream

- emulator
- sc

Use a sparse grid as a surrogate for the truth model when performing Bayesian inference.

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		sparse_grid_level	Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Description

Use a sparse grid as a surrogate for the truth model when performing Bayesian inference.

sparse_grid_level

- Keywords Area
- method
- bayes_calibration
- dream
- emulator
- sc
- sparse_grid_level

Set the maximum sparse grid level to be used when peforming sparse grid integration Cubature using spare grids

Specification

Alias: none

Argument(s): INTEGERLIST

Description

Multi-dimensional integration by the Smolyak sparse grid method (specified with sparse_grid_level and, optionally, dimension_preference). The underlying one-dimensional integration rules are the same as for the tensor-product quadrature case; however, the default rule selection is nested for sparse grids (Genz-Keister for normals/transformed normals and Gauss-Patterson for uniforms/transformed uniforms). This default can be overridden with an explicit non_nested specification (resulting in Gauss-Hermite for normals/transformed normals and Gauss-Legendre for uniforms/transformed uniforms). As for tensor quadrature, the dimension_preference specification enables the use of anisotropic sparse grids (refer to the PCE description in the User's Manual for the anisotropic index set constraint definition). Similar to anisotropic tensor grids, the dimension with greatest preference will have resolution at the full sparse_grid_level and all other dimension resolutions will be reduced in proportion to their reduced preference. For PCE with either isotropic or anisotropic sparse grids, a summation of tensor-product expansions is used, where each anisotropic tensor-product quadrature rule underlying the sparse grid construction results in its own anisotropic tensor-product expansion as described in case 1. These anisotropic tensor-product expansions are summed into a sparse PCE using the standard Smolyak summation (again, refer to the User's Manual for additional details). As for quadrature_order, the sparse_grid_level specification admits an array input for enabling specification of multiple grid resolutions used by certain advanced solution methodologies.

This keyword can be used when using sparse grid integration to calculate PCE coefficients or when generating samples for sparse grid collocation.

use_derivatives

- Keywords Area
- method
- bayes_calibration
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none

Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

likelihood_scale

- Keywords Area
- method
- bayes_calibration

- likelihood_scale

Scale the **log-likelihood** function

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): REAL

Description

Note that this changes the formulation of the calibration problem

calibrate_sigma

- Keywords Area
- method
- bayes_calibration
- calibrate_sigma

Calibrate the `sigma` term(s) in the likelihood function

Specification

Alias: none
Argument(s): none

Description

This feature allows one to calibrate the observational error (the `sigma` terms) as well as the model parameters.

samples

- Keywords Area
- method
- bayes_calibration
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- bayes_calibration
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- bayes_calibration
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.62 dace

- Keywords Area

- method

- dace

Design and Analysis of Computer Experiments

Topics

This keyword is related to the topics:

- package_ddace

- design_and_analysis_of_computer_experiments

- not_yet_reviewed

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
			grid	Grid Sampling
	Required*(Choose One)*	**Group 1**	random	Purely random samples
			oas	Orthogonal Array Sampling
			lhs	Use LHS method
			oa_lhs	Orthogonal Array Latin Hypercube Sampling
			box_behnken	Box-Behnken Design
			central_composite	Central Composite Design
	Optional		main_effects	ANOVA
	Optional		quality_metrics	Calculate metrics to assess the quality of the samples
	Optional		variance_based_-decomp	Global sensitivity analyis using variance-based decomposition
	Optional		fixed_seed	Reuse the same seed value for multiple sampling sets
	Optional		symbols	Number of replications in the sample set
	Optional		samples	Specify the number of samples
	Optional		seed	Specify the seed of the random number generator
	Optional		model_pointer	Specify which model block will be used by a method

Description

The Distributed Design and Analysis of Computer Experiments (DDACE) library provides the following DACE techniques:

1. grid sampling (`grid`)
2. pure random sampling (`random`)
3. orthogonal array sampling (`oas`)
4. latin hypercube sampling (`lhs`)
5. orthogonal array latin hypercube sampling (`oa_lhs`)
6. Box-Behnken (`box_behnken`)
7. central composite design (`central_composite`)

In addition to the selection of the method, there are keywords that affect the method outputs:

1. main_effects
2. quality_metrics
3. variance_based_decomp

And keywords that affect the sampling:

1. fixed_seed
2. symbols
3. samples
4. seed

See Also

These keywords may also be of interest:

- fsu_cvt
- fsu_quasi_mc
- psuade_moat

grid

- Keywords Area
- method
- dace
- grid

Grid Sampling

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

Grid Sampling

random

- Keywords Area
- method
- dace
- random

Purely random samples

Specification

Alias: none
Argument(s): none

Description

Purely random samples.

oas

- Keywords Area
- method
- dace
- oas

Orthogonal Array Sampling

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

Orthogonal Array Sampling

lhs

- Keywords Area
- method
- dace
- lhs

Use LHS method

Specification

Alias: none
Argument(s): none

Description

Use Latin Hypercube Sampling

oa_lhs

- Keywords Area
- method
- dace
- oa_lhs

Orthogonal Array Latin Hypercube Sampling

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

Orthogonal Array Latin Hypercube Sampling

box_behnken

- Keywords Area
- method
- dace
- box_behnken

Box-Behnken Design

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

Box-Behnken Design

central_composite

- Keywords Area
- method
- dace
- central_composite

Central Composite Design

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

Central Composite Design

main_effects

- Keywords Area
- method
- dace
- main_effects

ANOVA

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The `main_effects` control prints Analysis-of-Variance main effects results (e.g. ANOVA tables with p-values per variable). The `main_effects` control is only operational with the orthogonal arrays or Latin Hypercube designs, not for Box Behnken or Central Composite designs.

quality_metrics

- Keywords Area
- method
- dace
- quality_metrics

Calculate metrics to assess the quality of the samples

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

`quality_metrics` calculates four quality metrics relating to the volumetric spacing of the samples. The four quality metrics measure different aspects relating to the uniformity of point samples in hypercubes. Desirable properties of such point samples are:

- are the points equally spaced
- do the points cover the region
- and are they isotropically distributed
- with no directional bias in the spacing

The four quality metrics we report are:

- h: the point distribution norm, which is a measure of uniformity of the point distribution
- chi: a regularity measure, and provides a measure of local uniformity of a set of points
- tau: the second moment trace measure
- d: the second moment determinant measure

All of these values are scaled so that smaller is better (the smaller the metric, the better the uniformity of the point distribution).

Examples

Complete explanation of these measures can be found in [[38] "Gunzburger and Burkardt, 2004."].

variance_based_decomp

- Keywords Area
- method
- dace
- variance_based_decomp

Global sensitivity analyis using variance-based decomposition

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		drop_tolerance	Drop values lower than this tolerance

Description

Dakota can calculate sensitivity indices through variance based decomposition using the keyword `variance_based_decomp`. These indicate how important the uncertainty in each input variable is in contributing to the output variance.

`variance_based_decomp` is extremely computationally intensive since replicated sets of sample values are evaluated. If the user specified a number of samples, N, and a number of nondeterministic variables, M, variance-based decomposition requires the evaluation of N∗(M+2) samples. **Note that specifying this keyword will increase the number of function evaluations above the number requested with the `samples` keyword** .

Because of the computational cost, `variance_based_decomp` is turned off as a default.

To obtain sensitivity indices that are reasonably accurate, we recommend that N, the number of samples, be at least one hundred and preferably several hundred or thousands.

Theory

In this context, we take sensitivity analysis to be global, not local as when calculating derivatives of output variables with respect to input variables. Our definition is similar to that of [[74] "Saltelli et al., 2004"]: "The study of how uncertainty in the output of a model can be apportioned to different sources of uncertainty in the model input."

Variance based decomposition is a way of using sets of samples to understand how the variance of the output behaves, with respect to each input variable. A larger value of the sensitivity index, S_i, means that the uncertainty in the input variable i has a larger effect on the variance of the output. More details on the calculations and interpretation of the sensitivity indices can be found in [[74] "Saltelli et al., 2004"] and [[88] "Weirs et al., 2010"].

drop_tolerance

- Keywords Area
- method
- dace
- variance_based_decomp
- drop_tolerance

Drop values lower than this tolerance

Specification

Alias: none

Argument(s): REAL

Description

Drop values lower than this tolerance. For PCE which outputs main, interaction, and total effects by default, the univariate_effects option has been added to allow suppression of the interaction effects since the output volume of these results can be prohibitive for high dimensional problems. Similar to suppression of these interactions is the covariance control, which can be selected to be diagonal_covariance or full_covariance, with the former supporting suppression of the off-diagonal covariance terms (to save compute and memory resources and reduce output volume).

fixed_seed

- Keywords Area
- method
- dace
- fixed_seed

Reuse the same seed value for multiple sampling sets

Specification

Alias: none
Argument(s): none

Description

The `fixed_seed` flag is relevant if multiple sampling sets will be generated. This occurs when using advance methods (e.g., surrogate-based optimization, optimization under uncertainty).

Specifying this keyword results in the reuse of the same seed value for each of these multiple sampling sets, which can be important for reducing variability in the sampling results.

However, this behavior is not the default as the repetition of the same sampling pattern can result in a modeling weakness that an optimizer could potentially exploit (resulting in actual reliabilities that are lower than the estimated reliabilities). In either case (`fixed_seed` or not), the study is repeatable if the user specifies a `seed` and the study is random is the user omits a `seed` specification.

symbols

- Keywords Area
- method
- dace
- symbols

Number of replications in the sample set

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGER

Description

`symbols` is related to the number of replications in the sample set (a larger number of symbols equates to more stratification and fewer replications).

samples

- Keywords Area
- method
- dace
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- dace
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- dace
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.63 fsu_cvt

- Keywords Area
- method
- fsu_cvt

Design of Computer Experiments - Centroidal Voronoi Tessellation

Topics

This keyword is related to the topics:

- package_fsudace
- design_and_analysis_of_computer_experiments
- not_yet_reviewed

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		latinize	Adjust samples to improve the discrepancy of the marginal distributions
	Optional		quality_metrics	Calculate metrics to assess the quality of the samples
	Optional		variance_based_-decomp	Global sensitivity analyis using variance-based decomposition
	Optional		fixed_seed	Reuse the same seed value for multiple sampling sets
	Optional		trial_type	Specify how the trial samples are generated
	Optional		num_trials	The number of secondary sample points generated to adjust the location of the primary sample points
	Optional		samples	Specify the number of samples
	Optional		seed	Specify the seed of the random number generator
	Optional		model_pointer	Specify which model block will be used by a method

Description

The FSU Centroidal Voronoi Tessellation method (`fsu_cvt`) produces a set of sample points that are (approximately) a Centroidal Voronoi Tessellation. The primary feature of such a set of points is that they have good volumetric spacing; the points tend to arrange themselves in a pattern of cells that are roughly the same shape.

To produce this set of points, an almost arbitrary set of initial points is chosen, and then an internal set of iterations is carried out. These iterations repeatedly replace the current set of sample points by an estimate of the centroids of the corresponding Voronoi subregions. [[17] "Du, Faber, and Gunzburger, 1999"].

The user may generally ignore the details of this internal iteration. If control is desired, however, there are a few variables with which the user can influence the iteration. The user may specify:

- `max_iterations`, the number of iterations carried out (this is a method independent control, see max_-

iterations

- `num_trials`, the number of secondary sample points generated to adjust the location of the primary sample points
- `trial_type`, which controls how these secondary sample points are generated

This method generates sets of uniform random variables on the interval [0,1]. If the user specifies lower and upper bounds for a variable, the [0,1] samples are mapped to the [lower, upper] interval.

Theory

This method is designed to generate samples with the goal of low discrepancy. Discrepancy refers to the nonuniformity of the sample points within the hypercube.

Discrepancy is defined as the difference between the actual number and the expected number of points one would expect in a particular set B (such as a hyper-rectangle within the unit hypercube), maximized over all such sets. Low discrepancy sequences tend to cover the unit hypercube reasonably uniformly.

Centroidal Voronoi Tessellation does very well volumetrically: it spaces the points fairly equally throughout the space, so that the points cover the region and are isotropically distributed with no directional bias in the point placement. There are various measures of volumetric uniformity which take into account the distances between pairs of points, regularity measures, etc. Note that Centroidal Voronoi Tessellation does not produce low-discrepancy sequences in lower dimensions. The lower-dimension (such as 1-D) projections of Centroidal Voronoi Tessellation can have high discrepancy.

See Also

These keywords may also be of interest:

- dace
- fsu_quasi_mc
- psuade_moat

latinize

- Keywords Area
- method
- fsu_cvt
- latinize

Adjust samples to improve the discrepancy of the marginal distributions

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The `latinize` control takes the samples and "latinizes" them, meaning that each original sample is moved so that it falls into one strata or bin in each dimension as in Latin Hypercube sampling. The default setting is NOT to latinize. However, one may be interested in doing this in situations where one wants better discrepancy of the 1-dimensional projections (the marginal distributions).

quality_metrics

- Keywords Area
- method
- fsu_cvt
- quality_metrics

Calculate metrics to assess the quality of the samples

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

`quality_metrics` calculates four quality metrics relating to the volumetric spacing of the samples. The four quality metrics measure different aspects relating to the uniformity of point samples in hypercubes. Desirable properties of such point samples are:

- are the points equally spaced
- do the points cover the region
- and are they isotropically distributed
- with no directional bias in the spacing

The four quality metrics we report are:

- h: the point distribution norm, which is a measure of uniformity of the point distribution
- chi: a regularity measure, and provides a measure of local uniformity of a set of points

- tau: the second moment trace measure
- d: the second moment determinant measure

All of these values are scaled so that smaller is better (the smaller the metric, the better the uniformity of the point distribution).

Examples

Complete explanation of these measures can be found in [[38] "Gunzburger and Burkardt, 2004."].

variance_based_decomp

- Keywords Area
- method
- fsu_cvt
- variance_based_decomp

Global sensitivity analyis using variance-based decomposition

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		drop_tolerance	Drop values lower than this tolerance

Description

Dakota can calculate sensitivity indices through variance based decomposition using the keyword `variance_based_decomp`. These indicate how important the uncertainty in each input variable is in contributing to the output variance.

`variance_based_decomp` is extremely computationally intensive since replicated sets of sample values are evaluated. If the user specified a number of samples, N, and a number of nondeterministic variables, M, variance-based decomposition requires the evaluation of N∗(M+2) samples. **Note that specifying this keyword will increase the number of function evaluations above the number requested with the `samples` keyword** .

Because of the computational cost, `variance_based_decomp` is turned off as a default.

To obtain sensitivity indices that are reasonably accurate, we recommend that N, the number of samples, be at least one hundred and preferably several hundred or thousands.

Theory

In this context, we take sensitivity analysis to be global, not local as when calculating derivatives of output variables with respect to input variables. Our definition is similar to that of [[74] "Saltelli et al., 2004"]: "The study of how uncertainty in the output of a model can be apportioned to different sources of uncertainty in the model input."

Variance based decomposition is a way of using sets of samples to understand how the variance of the output behaves, with respect to each input variable. A larger value of the sensitivity index, S_i, means that the uncertainty in the input variable i has a larger effect on the variance of the output. More details on the calculations and interpretation of the sensitivity indices can be found in [[74] "Saltelli et al., 2004"] and [[88] "Weirs et al., 2010"].

drop_tolerance

- Keywords Area
- method
- fsu_cvt
- variance_based_decomp
- drop_tolerance

Drop values lower than this tolerance

Specification

Alias: none
Argument(s): REAL

Description

Drop values lower than this tolerance. For PCE which outputs main, interaction, and total effects by default, the univariate_effects option has been added to allow suppression of the interaction effects since the output volume of these results can be prohibitive for high dimensional problems. Similar to suppression of these interactions is the covariance control, which can be selected to be diagonal_covariance or full_covariance, with the former supporting suppression of the off-diagonal covariance terms (to save compute and memory resources and reduce output volume).

fixed_seed

- Keywords Area
- method
- fsu_cvt
- fixed_seed

Reuse the same seed value for multiple sampling sets

Specification

Alias: none
Argument(s): none

Description

The `fixed_seed` flag is relevant if multiple sampling sets will be generated. This occurs when using advance methods (e.g., surrogate-based optimization, optimization under uncertainty).

Specifying this keyword results in the reuse of the same seed value for each of these multiple sampling sets, which can be important for reducing variability in the sampling results.

However, this behavior is not the default as the repetition of the same sampling pattern can result in a modeling weakness that an optimizer could potentially exploit (resulting in actual reliabilities that are lower than the estimated reliabilities). In either case (`fixed_seed` or not), the study is repeatable if the user specifies a `seed` and the study is random is the user omits a `seed` specification.

trial_type

- Keywords Area
- method
- fsu_cvt
- trial_type

Specify how the trial samples are generated

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none

Argument(s): none

	Required/-Optional **Required***(Choose One)*	Description of Group **Group 1**	Dakota Keyword grid	Dakota Keyword Description Samples on a regular grid
			halton	Generate samples from a Halton sequence
			random	Purely random samples

Description

The user has the option to specify the method by which the trials are created to adjust the centroids. The `trial_type` can be one of three types:

- `random`, where points are generated randomly
- `halton`, where points are generated according to the Halton sequence
- `grid`, where points are placed on a regular grid over the hyperspace.

grid

- Keywords Area
- method
- fsu_cvt
- trial_type
- grid

Samples on a regular grid

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

Points are placed on a regular grid over the hyperspace.

halton

- Keywords Area
- method
- fsu_cvt
- trial_type
- halton

Generate samples from a Halton sequence

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The quasi-Monte Carlo sequences of Halton are deterministic sequences determined by a set of prime bases.

Generally, we recommend that the user leave the default setting for the bases, which are the lowest primes. Thus, if one wants to generate a sample set for 3 random variables, the default bases used are 2, 3, and 5 in the Halton sequence. To give an example of how these sequences look, the Halton sequence in base 2 starts with points 0.5, 0.25, 0.75, 0.125, 0.625, etc. The first few points in a Halton base 3 sequence are 0.33333, 0.66667, 0.11111, 0.44444, 0.77777, etc. Notice that the Halton sequence tends to alternate back and forth, generating a point closer to zero then a point closer to one. An individual sequence is based on a radix inverse function defined on a prime base. The prime base determines how quickly the [0,1] interval is filled in.

Theory

For more information about these sequences, see [[43] "Halton, 1960", [44] "Halton and Smith, 1964", and [1] "Kocis and Whiten, 1997"].

random

- Keywords Area
- method
- fsu_cvt
- trial_type
- random

Purely random samples

Specification

Alias: none

Argument(s): none

Description

Purely random samples.

num_trials

- Keywords Area
- method
- fsu_cvt
- num_trials

The number of secondary sample points generated to adjust the location of the primary sample points

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGER

Description

In general, the variable with the most influence on the quality of the final sample set is `num_trials`, which determines how well the Voronoi subregions are sampled.

Generally, `num_trials` should be "large", certainly much bigger than the number of sample points being requested; a reasonable value might be 10,000, but values of 100,000 or 1 million are not unusual.

samples

- Keywords Area
- method
- fsu_cvt
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- fsu_cvt
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- fsu_cvt
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.64 psuade_moat

- Keywords Area
- method
- psuade_moat

Morris One-at-a-Time

Topics

This keyword is related to the topics:

- package_psuade
- design_and_analysis_of_computer_experiments
- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		partitions	Number of partitions of each variable
	Optional		samples	Specify the number of samples
	Optional		seed	Specify the seed of the random number generator
	Optional		model_pointer	Specify which model block will be used by a method

Description

!!! Check on Latex

The Morris One-At-A-Time (MOAT) method, originally proposed by Morris [[63] "M.D. Morris, 1991"], is a screening method, designed to explore a computational model to distinguish between input variables that have negligible, linear and additive, or nonlinear or interaction effects on the output. The computer experiments performed consist of individually randomized designs which vary one input factor at a time to create a sample of its elementary effects.

The number of samples (`samples`) must be a positive integer multiple of (number of continuous design variable + 1) and will be automatically adjusted if misspecified.

The number of partitions (`partitions`) applies to each variable being studied and must be odd (the number of MOAT levels per variable is partitions + 1). This will also be adjusted at runtime as necessary.

For information on practical use of this method, see [[74] "Saltelli, et al., 2004"].

Theory

With MOAT, each dimension of a $k-$dimensional input space is uniformly partitioned into p levels, creating a grid of p^k points $\mathbf{x} \in \mathbf{R^k}$ at which evaluations of the model $y(\mathbf{x})$ might take place. An elementary effect corresponding to input i is computed by a forward difference

$$d_i(\mathbf{x}) = \frac{y(\mathbf{x} + \Delta \mathbf{e}_i) - y(\mathbf{x})}{\Delta},$$

where e_i is the ith coordinate vector, and the step Δ is typically taken to be large (this is not intended to be a local derivative approximation). In the present implementation of MOAT, for an input variable scaled to $[0, 1]$, $\Delta = \frac{p}{2(p-1)}$, so the step used to find elementary effects is slightly larger than half the input range.

The distribution of elementary effects d_i over the input space characterizes the effect of input i on the output of interest. After generating r samples from this distribution, their mean,

$$\mu_i = \frac{1}{r} \sum_{j=1}^{r} d_i^{(j)}$$

modified mean

$$\mu_i^* = \frac{1}{r}\sum_{j=1}^{r} |d_i^{(j)}|,$$

(using absolute value) and standard deviation

$$\sigma_i = \sqrt{\frac{1}{r}\sum_{j=1}^{r}\left(d_i^{(j)} - \mu_i\right)^2}$$

are computed for each input i. The mean and modified mean give an indication of the overall effect of an input on the output. Standard deviation indicates nonlinear effects or interactions, since it is an indicator of elementary effects varying throughout the input space.

partitions

- Keywords Area
- method
- psuade_moat
- partitions

Number of partitions of each variable

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGERLIST

Description

Described on the parent page.

samples

- Keywords Area
- method
- psuade_moat
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

seed

- Keywords Area
- method
- psuade_moat
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

model_pointer

- Keywords Area
- method
- psuade_moat
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.65 local_evidence

- Keywords Area
- method
- local_evidence

Evidence theory with evidence measures computed with local optimization methods

Topics

This keyword is related to the topics:

- epistemic_uncertainty_quantification_methods
- evidence_theory

Specification

Alias: nond_local_evidence

Argument(s): none

	Required/-Optional **Optional***(Choose One)*	Description of Group **Group 1**	Dakota Keyword sqp	Dakota Keyword Description Use a SQP method for optimization
			nip	Use a nonlinear interior point method for optimization
	Optional		response_levels	Specify the response values at which to estimate relevant statistics
	Optional		probability_levels	Specify probability levels at which to estimate the corresponding response value

	Optional	gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional	distribution	Select between cumulative or complementary-cumulative functions
	Optional	model_pointer	Specify which model block will be used by a method

Description

See the topic page evidence_theory for important background information and usage notes

Two local optimization methods are available: `sqp` (sequential quadratic programming or `nip` (nonlinear interior point method).

See Also

These keywords may also be of interest:

- global_evidence
- global_interval_est
- local_interval_est

sqp

- Keywords Area
- method
- local_evidence
- sqp

Use a SQP method for optimization

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

Description

Which one???

nip

- Keywords Area
- method
- local_evidence
- nip

Use a nonlinear interior point method for optimization

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

Description

Which one???

response_levels

- Keywords Area
- method
- local_evidence
- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_-levels` correspond to which response
	Optional		compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_-response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- local_evidence
- response_levels
- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- local_evidence
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** probabilities	**Dakota Keyword Description** Compute probabilities
			gen_reliabilities	Compute generalized reliabilities

	Optional	system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- local_evidence
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- local_evidence
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- local_evidence
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	series	?????
			parallel	?????

series

- Keywords Area
- method
- local_evidence
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

parallel

- Keywords Area
- method
- local_evidence
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

probability_levels

- Keywords Area
- method
- local_evidence
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- local_evidence
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none

Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- local_evidence
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- local_evidence
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

distribution

- Keywords Area
- method
- local_evidence
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** cumulative	**Dakota Keyword Description** Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.
This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- local_evidence
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- local_evidence
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

model_pointer

- Keywords Area
- method
- local_evidence
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.66 local_interval_est

- Keywords Area
- method
- local_interval_est

Interval analysis using local optimization

Topics

This keyword is related to the topics:

- uncertainty_quantification
- epistemic_uncertainty_quantification_methods
- interval_estimation

Specification

Alias: nond_local_interval_est
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	sqp	Use a SQP method for optimization
			nip	Use a nonlinear interior point method for optimization
	Optional		model_pointer	Specify which model block will be used by a method

Description

Interval analysis using local methods (`local_interval_est`). If the problem is amenable to local optimization methods (e.g. can provide derivatives or use finite difference method to calculate derivatives), then one can use one of two local methods to calculate these bounds.

- `sqp`
- `nip`

Theory

In interval analysis, one assumes that nothing is known about an epistemic uncertain variable except that its value lies somewhere within an interval. In this situation, it is NOT assumed that the value has a uniform probability of occuring within the interval. Instead, the interpretation is that any value within the interval is a possible value or a potential realization of that variable. In interval analysis, the uncertainty quantification problem is one of determining the resulting bounds on the output (defining the output interval) given interval bounds on the inputs. Again, any output response that falls within the output interval is a possible output with no frequency information assigned to it.

See Also

These keywords may also be of interest:

- global_evidence
- global_interval_est
- local_evidence

sqp

- Keywords Area
- method
- local_interval_est
- sqp

Use a SQP method for optimization

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

Description

Which one???

nip

- Keywords Area
- method
- local_interval_est
- nip

Use a nonlinear interior point method for optimization

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

Description

Which one???

model_pointer

- Keywords Area
- method
- local_interval_est
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.67 local_reliability

- Keywords Area
- method
- local_reliability

Local reliability method

Topics

This keyword is related to the topics:

- uncertainty_quantification
- reliability_methods

Specification

Alias: nond_local_reliability
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		mpp_search	Specify which MPP search option to use
	Optional		response_levels	Specify the response values at which to estimate relevant statistics

	Optional	reliability_levels	Specify reliability levels at which the response values will be estimated
	Optional	distribution	Select between cumulative or complementary-cumulative functions
	Optional	probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional	gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional	model_pointer	Specify which model block will be used by a method

Description

Local reliability methods compute approximate response function distribution statistics based on specified uncertain variable probability distributions. Each of the local reliability methods can compute forward and inverse mappings involving response, probability, reliability, and generalized reliability levels.

The forward reliability analysis algorithm of computing reliabilities/probabilities for specified response levels is called the Reliability Index Approach (RIA), and the inverse reliability analysis algorithm of computing response levels for specified probability levels is called the Performance Measure Approach (PMA).

The different RIA/PMA algorithm options are specified using the `mpp_search` specification which selects among different limit state approximations that can be used to reduce computational expense during the MPP searches.

Theory

The Mean Value method (MV, also known as MVFOSM in [[42] "Haldar and Mahadevan, 2000"]) is the simplest, least-expensive method in that it estimates the response means, response standard deviations, and all CDF/CCDF forward/inverse mappings from a single evaluation of response functions and gradients at the uncertain variable means. This approximation can have acceptable accuracy when the response functions are nearly linear and their distributions are approximately Gaussian, but can have poor accuracy in other situations.

All other reliability methods perform an internal nonlinear optimization to compute a most probable point (MPP) of failure. A sign convention and the distance of the MPP from the origin in the transformed standard normal space ("u-space") define the reliability index, as explained in the section on Reliability Methods in the Uncertainty Quantification chapter of the Users Manual [[4] "Adams et al., 2010"]. The reliability can then be converted to a probability using either first- or second-order integration, may then be refined using importance sampling, and finally may be converted to a generalized reliability index.

See Also

These keywords may also be of interest:

- adaptive_sampling
- gpais
- global_reliability
- sampling
- importance_sampling
- polynomial_chaos
- stoch_collocation

mpp_search

- Keywords Area
- method
- local_reliability
- mpp_search

Specify which MPP search option to use

Topics

This keyword is related to the topics:

- uncertainty_quantification
- reliability_methods

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword x_taylor_mean	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	u_taylor_mean	
			x_taylor_mpp	
			u_taylor_mpp	
			x_two_point	
			u_two_point	

			no_approx	
	Optional*(Choose One)*	**Group 2**	sqp	Use a SQP method for optimization
			nip	Use a nonlinear interior point method for optimization
	Optional		integration	

Description

The `x_taylor_mean` MPP search option performs a single Taylor series approximation in the space of the original uncertain variables ("x-space") centered at the uncertain variable means, searches for the MPP for each response/probability level using this approximation, and performs a validation response evaluation at each predicted MPP. This option is commonly known as the Advanced Mean Value (AMV) method. The `u_taylor_mean` option is identical to the `x_taylor_mean` option, except that the approximation is performed in u-space. The `x_taylor_mpp` approach starts with an x-space Taylor series at the uncertain variable means, but iteratively updates the Taylor series approximation at each MPP prediction until the MPP converges. This option is commonly known as the AMV+ method. The `u_taylor_mpp` option is identical to the `x_taylor_mpp` option, except that all approximations are performed in u-space. The order of the Taylor-series approximation is determined by the corresponding `responses` specification and may be first or second-order. If second-order (methods named AMV^2 and AMV^2+ in [[21] "Eldred and Bichon, 2006"]), the series may employ analytic, finite difference, or quasi Hessians (BFGS or SR1). The `x_two_point` MPP search option uses an x-space Taylor series approximation at the uncertain variable means for the initial MPP prediction, then utilizes the Two-point Adaptive Nonlinear Approximation (TANA) outlined in [[92] "Xu and Grandhi, 1998"] for all subsequent MPP predictions. The `u_two_point` approach is identical to `x_two_point`, but all the approximations are performed in u-space. The `x_taylor_mpp` and `u_taylor_mpp`, `x_two_point` and `u_two_point` approaches utilize the `max_iterations` and `convergence_tolerance` method independent controls to control the convergence of the MPP iterations (the maximum number of MPP iterations per level is limited by `max_iterations`, and the MPP iterations are considered converged when $\| \mathbf{u}^{(k+1)} - \mathbf{u}^{(k)} \|_2 <$ `convergence_tolerance`). And, finally, the `no_approx` option performs the MPP search on the original response functions without the use of any approximations. The optimization algorithm used to perform these MPP searches can be selected to be either sequential quadratic programming (uses the `npsol_sqp` optimizer) or nonlinear interior point (uses the `optpp_q_newton` optimizer) algorithms using the `sqp` or `nip` keywords.

In addition to the MPP search specifications, one may select among different integration approaches for computing probabilities at the MPP by using the `integration` keyword followed by either `first_order` or `second_order`. Second-order integration employs the formulation of [[50] "Hohenbichler and Rackwitz, 1988"] (the approach of [[12] "Breitung, 1984"] and the correction of [[51] "Hong 1999"] are also implemented, but are not active). Combining the `no_approx` option of the MPP search with first- and second-order integrations results in the traditional first- and second-order reliability methods (FORM and SORM). These integration approximations may be subsequently refined using importance sampling. The `refinement` specification allows the seletion of basic importance sampling (`import`), adaptive importance sampling (`adapt_import`), or multimodal adaptive importance sampling (`mm_adapt_import`), along with the specification of number of samples (`samples`) and random seed (`seed`). Additional details on these methods are available in [[23] "Eldred et al., 2004"] and [[21] "Eldred and Bichon, 2006"] and in the Uncertainty Quantification Capabilities chapter of the Users Manual [[4] "Adams et al., 2010"].

x_taylor_mean

- Keywords Area

- method
- local_reliability
- mpp_search
- x_taylor_mean

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: none
Argument(s): none

Description

This mpp_search option performs a single Taylor series approximation in the space of the original uncertain variables (”x-space”) centered at the uncertain variable means, searches for the MPP for each response/probability level using this approximation, and performs a validation response evaluation at each predicted MPP. This option is commonly known as the Advanced Mean Value (AMV) method.

u_taylor_mean

- Keywords Area
- method
- local_reliability
- mpp_search
- u_taylor_mean

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: none
Argument(s): none

Description

This mpp_search option performs a single Taylor series approximation in the transformed space of the uncertain variables ("u-space") centered at the uncertain variable means. This option is commonly known as the Advanced Mean Value (AMV) method, but is performed in u-space instead of x-space.

x_taylor_mpp

- Keywords Area
- method
- local_reliability
- mpp_search
- x_taylor_mpp

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: none
Argument(s): none

Description

This mpp_search option starts with an x-space Taylor series at the uncertain variable means, but iteratively updates the Taylor series approximation at each MPP prediction until the MPP converges. This option is commonly known as the AMV+ method.

u_taylor_mpp

- Keywords Area
- method
- local_reliability
- mpp_search
- u_taylor_mpp

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: none
Argument(s): none

Description

This mpp_search option starts with a u-space Taylor series at the uncertain variable means, and iteratively updates the Taylor series approximation at each MPP prediction until the MPP converges. This option is commonly known as the AMV+ method and is identify to x_taylor_mpp except that it is performed in u-space.

x_two_point

- Keywords Area
- method
- local_reliability
- mpp_search
- x_two_point

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: none
Argument(s): none

Description

This mpp_search option uses an x-space Taylor series approximation at the uncertain variable means for the initial MPP prediction, then utilizes the Two-point Adaptive Nonlinear Approximation (TANA) outlined in [Xu98 "Xu and Grandhi, 1998"] for all subsequent MPP predictions.

u_two_point

- Keywords Area
- method
- local_reliability
- mpp_search
- u_two_point

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: none
Argument(s): none

Description

This mpp_search option is identical to x_two_point, but it performs the Two-point Adaptive Nonlinear Approximation (TANA) in u-space instead of x-space.

no_approx

- Keywords Area
- method
- local_reliability
- mpp_search
- no_approx

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: none
Argument(s): none

Description

This mpp_search option performs the MPP search on the original response functions without the use of any approximations. Note that the use of the no_approx MPP search with first-order probability integration results in the traditional reliability method called FORM (First-Order Reliability Method). Similarly, the use of no_approx with second-order probability integration results in SORM (Second-Order Reliability Method).

sqp

- Keywords Area
- method
- local_reliability
- mpp_search
- sqp

Use a SQP method for optimization

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

Description

Which one???

nip

- Keywords Area
- method
- local_reliability
- mpp_search
- nip

Use a nonlinear interior point method for optimization

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

Description

Which one???

integration

- Keywords Area
- method
- local_reliability
- mpp_search
- integration

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	first_order	
			second_order	
	Optional		probability_-refinement	Allow refinement of probability and generalized reliability results using importance sampling

Description

This keyword controls how the probabilities at the MPP are computed: integration is followed by either first_order or second_order, indicating the order of the probability integration.

first_order

- Keywords Area
- method
- local_reliability
- mpp_search
- integration
- first_order

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: none
Argument(s): none

Description

First-order integration in local reliability methods uses the minimum Euclidean distance from the origin to the most probable point (MPP) in transformed space to compute the probability of failure. This distance, commonly called the reliability index Beta, is used to calculate the probability of failure by calculating the standard normal cumulative distribution function at -Beta.

second_order

- Keywords Area
- method
- local_reliability
- mpp_search
- integration
- second_order

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: none
Argument(s): none

Description

Second-order integration in local reliability methods modifies the first-order integration approach to apply a curvature correction. This correction is based on the formulation of [Hoh88 "Hohenbichler and Rackwitz, 1988"].

probability_refinement

- Keywords Area
- method
- local_reliability
- mpp_search
- integration
- probability_refinement

Allow refinement of probability and generalized reliability results using importance sampling

Topics

This keyword is related to the topics:

- problem

Specification

Alias: sample_refinement
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	import	Sampling option
			adapt_import	Importance sampling option
			mm_adapt_import	Sampling option
	Optional		refinement_-samples	Specify the number of samples used to improve a probabilty estimate.

	Optional	seed	Specify the seed of the random number generator

Description

Allow refinement of probability and generalized reliability results using importance sampling

The `probability_refinement` density reweighting must account for additional non-Gaussian cases. This option allows for refinement of probability and generalized reliability results using importance sampling.

import

- Keywords Area
- method
- local_reliability
- mpp_search
- integration
- probability_refinement
- import

Sampling option

Specification

Alias: none

Argument(s): none

Description

`import` centers a sampling density at one of the initial LHS samples identified in the failure region. It then generates the importance samples, weights them by their probability of occurence given the original density, and calculates the required probability (CDF or CCDF level).

adapt_import

- Keywords Area
- method
- local_reliability
- mpp_search
- integration
- probability_refinement
- adapt_import

Importance sampling option

Specification

Alias: none
Argument(s): none

Description

`adapt_import` centers a sampling density at one of the initial LHS samples identified in the failure region. It then generates the importance samples, weights them by their probability of occurence given the original density, and calculates the required probability (CDF or CCDF level). This continues iteratively until the failure probability estimate converges.

mm_adapt_import

- Keywords Area
- method
- local_reliability
- mpp_search
- integration
- probability_refinement
- mm_adapt_import

Sampling option

Specification

Alias: none
Argument(s): none

Description

`mm_adapt_import` starts with all of the samples located in the failure region to build a multimodal sampling density. First, it uses a small number of samples around each of the initial samples in the failure region. Note that these samples are allocated to the different points based on their relative probabilities of occurrence: more probable points get more samples. This early part of the approach is done to search for "representative" points. Once these are located, the multimodal sampling density is set and then `mm_adapt_import` proceeds similarly to `adapt_import` (sample until convergence).

refinement_samples

- Keywords Area
- method
- local_reliability
- mpp_search

- integration
- probability_refinement
- refinement_samples

Specify the number of samples used to improve a probabilty estimate.

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples used to improve a probabilty estimate. If using uni-modal sampling all samples are assigned to the sampling center. If using multi-modal sampling the samples are split between mutiple samples according to some internally computed weights.

seed

- Keywords Area
- method
- local_reliability
- mpp_search
- integration
- probability_refinement
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

response_levels

- Keywords Area
- method
- local_reliability
- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_levels` correspond to which response
	Optional		compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- local_reliability
- response_levels
- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- local_reliability
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none
Argument(s): none

	Required/-Optional	**Description of Group**	**Dakota Keyword**	**Dakota Keyword Description**
	Required*(Choose One)*	**Group 1**	probabilities	Compute probabilities
			reliabilities	Compute reliabilities
			gen_reliabilities	Compute generalized reliabilities

	Optional	system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- local_reliability
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

reliabilities

- Keywords Area
- method
- local_reliability
- response_levels
- compute
- reliabilities

Compute reliabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- local_reliability
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none
Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- local_reliability
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	series	?????
			parallel	?????

series

- Keywords Area
- method
- local_reliability
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

parallel

- Keywords Area
- method
- local_reliability
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

reliability_levels

- Keywords Area
- method
- local_reliability
- reliability_levels

Specify reliability levels at which the response values will be estimated

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_reliability_-levels	Specify which `reliability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF reliabilities by projecting out the prescribed number of sample standard deviations from the sample mean.

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_reliability_levels

- Keywords Area
- method
- local_reliability
- reliability_levels
- num_reliability_levels

Specify which `reliability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

distribution

- Keywords Area
- method
- local_reliability
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	cumulative	Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.
This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- local_reliability
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- local_reliability
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- local_reliability
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- local_reliability
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- local_reliability
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- local_reliability
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none

Argument(s): INTEGERLIST

Description

See parent page

model_pointer

- Keywords Area
- method
- local_reliability
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.68 global_reliability

- Keywords Area
- method
- global_reliability

Global reliability methods

Topics

This keyword is related to the topics:

- uncertainty_quantification
- reliability_methods
- problem

Specification

Alias: nond_global_reliability
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	x_gaussian_process	Create GP surrogate in x-space
			u_gaussian_process	Create GP surrogate in u-space
	Optional*(Choose One)*	**Group 2**	surfpack	Use the Surfpack version of Gaussian Process surrogates
			dakota	Select the built in Gaussian Process surrogate
	Optional		import_points_file	Specify file containing points (variable and response values)
	Optional		export_points_file	Specify file to write out the responses from the surrogate
	Optional		use_derivatives	Use derivative data to construct surrogate models
	Optional		seed	Specify the seed of the random number generator
	Optional		rng	Choose a random number generator
	Optional		response_levels	Specify the response values at which to estimate relevant statistics
	Optional		distribution	Select between cumulative or complementary-cumulative functions

	Optional	probability_levels	Specify probability levels at which to estimate the corresponding response value
	Optional	gen_reliability_-levels	Specify generalized relability levels at which to estimate the corresponding response value
	Optional	model_pointer	Specify which model block will be used by a method

Description

These methods do not support forward/inverse mappings involving `reliability_levels`, since they never form a reliability index based on distance in u-space. Rather they use a Gaussian process model to form an approximation to the limit state (based either in x-space via the `x_gaussian_process` specification or in u-space via the `u_gaussian_process` specification), followed by probability estimation based on multimodal adaptive importance sampling (see [[10] "Bichon et al., 2007"]) and [[11] "Bichon et al., 2008"]). These probability estimates may then be transformed into generalized reliability levels if desired. At this time, inverse reliability analysis (mapping probability or generalized reliability levels into response levels) is not implemented.

The Gaussian process model approximation to the limit state is formed over the aleatory uncertain variables by default, but may be extended to also capture the effect of design, epistemic uncertain, and state variables. If this is desired, one must use the appropriate controls to specify the active variables in the variables specification block.

See Also

These keywords may also be of interest:

- adaptive_sampling
- gpais
- local_reliability
- sampling
- importance_sampling
- polynomial_chaos
- stoch_collocation

x_gaussian_process

- Keywords Area
- method

- global_reliability
- x_gaussian_process

Create GP surrogate in x-space

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: x_kriging
Argument(s): none

u_gaussian_process

- Keywords Area
- method
- global_reliability
- u_gaussian_process

Create GP surrogate in u-space

Topics

This keyword is related to the topics:

- reliability_methods

Specification

Alias: u_kriging
Argument(s): none

surfpack

- Keywords Area
- method
- global_reliability
- surfpack

Use the Surfpack version of Gaussian Process surrogates

Specification

Alias: none
Argument(s): none

Description

This keyword specifies the use of the Gaussian process that is incorporated in our surface fitting library called Surfpack.

Several user options are available:

1. Optimization methods:

   Maximum Likelihood Estimation (MLE) is used to find the optimal values of the hyper-parameters governing the trend and correlation functions. By default the global optimization method DIRECT is used for MLE, but other options for the optimization method are available. See optimization_method.

   The total number of evaluations of the likelihood function can be controlled using the `max_trials` keyword followed by a positive integer. Note that the likelihood function does not require running the "truth" model, and is relatively inexpensive to compute.

2. Trend Function:

   The GP models incorporate a parametric trend function whose purpose is to capture large-scale variations. See trend.

3. Correlation Lengths:

   Correlation lengths are usually optimized by Surfpack, however, the user can specify the lengths manually. See correlation_lengths.

4. Ill-conditioning

   One of the major problems in determining the governing values for a Gaussian process or Kriging model is the fact that the correlation matrix can easily become ill-conditioned when there are too many input points close together. Since the predictions from the Gaussian process model involve inverting the correlation matrix, ill-conditioning can lead to poor predictive capability and should be avoided.

   Note that a sufficiently bad sample design could require correlation lengths to be so short that any interpolatory GP model would become inept at extrapolation and interpolation.

   The `surfpack` model handles ill-conditioning internally by default, but behavior can be modified using

5. Gradient Enhanced Kriging (GEK).

   The `use_derivatives` keyword will cause the Surfpack GP to be constructed from a combination of function value and gradient information (if available).

   See notes in the Theory section.

Theory

Gradient Enhanced Kriging

Incorporating gradient information will only be beneficial if accurate and inexpensive derivative information is available, and the derivatives are not infinite or nearly so. Here "inexpensive" means that the cost of evaluating a function value plus gradient is comparable to the cost of evaluating only the function value, for example gradients computed by analytical, automatic differentiation, or continuous adjoint techniques. It is not cost effective to use derivatives computed by finite differences. In tests, GEK models built from finite difference derivatives were also

significantly less accurate than those built from analytical derivatives. Note that GEK's correlation matrix tends to have a significantly worse condition number than Kriging for the same sample design.

This issue was addressed by using a pivoted Cholesky factorization of Kriging's correlation matrix (which is a small sub-matrix within GEK's correlation matrix) to rank points by how much unique information they contain. This reordering is then applied to whole points (the function value at a point immediately followed by gradient information at the same point) in GEK's correlation matrix. A standard non-pivoted Cholesky is then applied to the reordered GEK correlation matrix and a bisection search is used to find the last equation that meets the constraint on the (estimate of) condition number. The cost of performing pivoted Cholesky on Kriging's correlation matrix is usually negligible compared to the cost of the non-pivoted Cholesky factorization of GEK's correlation matrix. In tests, it also resulted in more accurate GEK models than when pivoted Cholesky or whole-point-block pivoted Cholesky was performed on GEK's correlation matrix.

?? This correlation function results in a response surface that is C^{∞}-continuous.

dakota

- Keywords Area
- method
- global_reliability
- dakota

Select the built in Gaussian Process surrogate

Specification

Alias: none

Argument(s): none

Description

A second version of GP surrogates was available in prior versions of Dakota. **For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.**

Historically these models were drastically different, but in Dakota 5.1, they became quite similar. They now differ in that the Surfpack GP has a richer set of features/options and tends to be more accurate than the Dakota version. Due to how the Surfpack GP handles ill-conditioned correlation matrices (which significantly contributes to its greater accuracy), the `Surfpack` GP can be a factor of two or three slower than Dakota's. As of Dakota 5.2, the Surfpack implementation is the default in all contexts except Bayesian calibration.

More details on the `gaussian_process` dakota model can be found in[59].

Dakota's GP deals with ill-conditioning in two ways. First, when it encounters a non-invertible correlation matrix it iteratively increases the size of a "nugget," but in such cases the resulting approximation smooths rather than interpolates the data. Second, it has a `point_selection` option (default off) that uses a greedy algorithm to select a well-spaced subset of points prior to the construction of the GP. In this case, the GP will only interpolate the selected subset. Typically, one should not need point selection in trust-region methods because a small number of points are used to develop a surrogate within each trust region. Point selection is most beneficial when constructing with a large number of points, typically more than order one hundred, though this depends on the number of variables and spacing of the sample points.

This differs from the `point_selection` option of the Dakota GP which initially chooses a well-spaced subset of points and finds the correlation parameters that are most likely for that one subset.

import_points_file

- Keywords Area
- method
- global_reliability
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none
Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- global_reliability
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- global_reliability
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- method
- global_reliability
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none

Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional(*Choose One*)	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- method
- global_reliability

- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- global_reliability
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

use_derivatives

- Keywords Area
- method
- global_reliability
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none
Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

seed

- Keywords Area
- method
- global_reliability
- seed

Specify the seed of the random number generator

Specification

Alias: none
Argument(s): INTEGER

Description

The random `seed` control provides a mechanism for making a stochastic optimization repeatable. That is, the use of the same random seed in identical studies will generate identical results.

rng

- Keywords Area
- method
- global_reliability
- rng

Choose a random number generator

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	mt19937	Use the Mersenne twister random number generator
			rnum2	Random Number Generator 2

Description

There are currently two options.

mt19937

- Keywords Area
- method
- global_reliability
- rng
- mt19937

Use the Mersenne twister random number generator

Specification

Alias: none
Argument(s): none

Description

See wikipedia: http://en.wikipedia.org/wiki/Mersenne_twister

rnum2

- Keywords Area
- method
- global_reliability
- rng
- rnum2

Random Number Generator 2

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

It's not the default for a reason...

response_levels

- Keywords Area
- method
- global_reliability
- response_levels

Specify the response values at which to estimate relevant statistics

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_response_-levels	Specify which `response_-levels` correspond to which response
	Optional		compute	Choose which forward mapping to compute at each response level

Description

The `response_levels` specification provides the target response values for which to compute probabilities, reliabilities, or generalized reliabilities (forward mapping).

The particular statistics depend on the method, and they include:

1. Reliabilities
2. CDF probabilities
3. CCDF probabilities

`num_response_levels` is used to specify which arguments of the `response_level` correspond to which response.

Examples

For example, specifying a `response_level` of 52.3 followed with `compute probabilities` will result in the calculation of the probability that the response value is less than or equal to 52.3, given the uncertain distributions on the inputs.

For an example with multiple responses, the following specification

```
response_levels = 1. 2. .1 .2 .3 .4 10. 20. 30.
  num_response_levels = 2 4 3
```

would assign the first two response levels (1., 2.) to response function 1, the next four response levels (.1, .2, .3, .4) to response function 2, and the final three response levels (10., 20., 30.) to response function 3. If the `num_-response_levels` key were omitted from this example, then the response levels would be evenly distributed among the response functions (three levels each in this case).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

A forward mapping involves computing the belief and plausibility probability level for a specified response level.

num_response_levels

- Keywords Area
- method
- global_reliability
- response_levels
- num_response_levels

Specify which `response_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

compute

- Keywords Area
- method
- global_reliability
- response_levels
- compute

Choose which forward mapping to compute at each response level

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** probabilities	**Dakota Keyword Description** Compute probabilities
			gen_reliabilities	Compute generalized reliabilities

	Optional	system	?????

Description

`compute` is used to select which forward mapping is calculated at each response level.

CDF/CCDF probabilities are calculated for specified response levels using a simple binning approach.

CDF/CCDF reliabilities are calculated for specified response levels by computing the number of sample standard deviations separating the sample mean from the response level.

probabilities

- Keywords Area
- method
- global_reliability
- response_levels
- compute
- probabilities

Compute probabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

gen_reliabilities

- Keywords Area
- method
- global_reliability
- response_levels
- compute
- gen_reliabilities

Compute generalized reliabilities

Specification

Alias: none

Argument(s): none

Description

See parent page.

system

- Keywords Area
- method
- global_reliability
- response_levels
- compute
- system

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	series	?????
			parallel	?????

series

- Keywords Area
- method
- global_reliability
- response_levels
- compute
- system
- series

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

parallel

- Keywords Area
- method
- global_reliability
- response_levels
- compute
- system
- parallel

?????

Topics

This keyword is related to the topics:

- not_yet_reviewed
- problem

Specification

Alias: none
Argument(s): none

distribution

- Keywords Area
- method
- global_reliability
- distribution

Select between cumulative or complementary-cumulative functions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	cumulative	Choose cumulative functions
			complementary	Choose complementary cumulative functions

Description

The selection of cumulative:

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

vs. complementary cumulative:

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

can be performed with the `distribution` keyword and the choice of `cumulative` or `complementary`.
This selection also defines the sign of the reliability or generalized reliability indices.

cumulative

- Keywords Area
- method
- global_reliability
- distribution
- cumulative

Choose cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CDF - cumulative distribution function
- CBF - cumulative belief function
- CPF - cumulative plausibility function

complementary

- Keywords Area
- method
- global_reliability
- distribution
- complementary

Choose complementary cumulative functions

Specification

Alias: none
Argument(s): none

Description

Applies to

- CCDF - complementary cumulative distribution function
- CCBF - complementary cumulative belief function
- CCPF - complementary cumulative plausibility function

probability_levels

- Keywords Area
- method
- global_reliability
- probability_levels

Specify probability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_probability_-levels	Specify which `probability_-levels` correspond to which response

Description

Response levels are calculated for specified CDF/CCDF probabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_probability_levels

- Keywords Area
- method
- global_reliability
- probability_levels
- num_probability_levels

Specify which `probability_levels` correspond to which response

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See parent page

gen_reliability_levels

- Keywords Area
- method
- global_reliability
- gen_reliability_levels

Specify generalized relability levels at which to estimate the corresponding response value

Specification

Alias: none
Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_gen_-reliability_levels	Specify which `gen_-reliability_-levels` correspond to which response

Description

Response levels are calculated for specified generalized reliabilities by indexing into a sorted samples array (the response levels computed are not interpolated and will correspond to one of the sampled values).

Theory

Sets of response-probability pairs computed with the forward/inverse mappings define either a cumulative distribution function (CDF) or a complementary cumulative distribution function (CCDF) for each response function.

In the case of evidence-based epistemic methods, this is generalized to define either cumulative belief and plausibility functions (CBF and CPF) or complementary cumulative belief and plausibility functions (CCBF and CCPF) for each response function.

An inverse mapping involves computing the belief and plausibility response level for either a specified probability level or a specified generalized reliability level (two results for each level mapping in the evidence-based epistemic case, instead of the one result for each level mapping in the aleatory case).

num_gen_reliability_levels

- Keywords Area
- method
- global_reliability
- gen_reliability_levels
- num_gen_reliability_levels

Specify which `gen_reliability_levels` correspond to which response

Specification

Alias: none

Argument(s): INTEGERLIST

Description

See parent page

model_pointer

- Keywords Area
- method
- global_reliability
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.69 fsu_quasi_mc

- Keywords Area
- method
- fsu_quasi_mc

Design of Computer Experiments - Quasi-Monte Carlo sampling

Topics

This keyword is related to the topics:

- package_fsudace
- design_and_analysis_of_computer_experiments
- not_yet_reviewed

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	halton	Generate samples from a Halton sequence
			hammersley	Use Hammersley sequences
	Optional		latinize	Adjust samples to improve the discrepancy of the marginal distributions
	Optional		quality_metrics	Calculate metrics to assess the quality of the samples
	Optional		variance_based_-decomp	Global sensitivity analyis using variance-based decomposition
	Optional		samples	Specify the number of samples
	Optional		fixed_sequence	Reuse the same sequence and samples for multiple sampling sets
	Optional		sequence_start	Choose where to start sampling the sequence
	Optional		sequence_leap	Specify how often the sequence is sampled
	Optional		prime_base	The prime numbers used to generate the sequence
	Optional		model_pointer	Specify which model block will be used by a method

Description

Quasi-Monte Carlo methods produce low discrepancy sequences, especially if one is interested in the uniformity of projections of the point sets onto lower dimensional faces of the hypercube (usually 1-D: how well do the marginal distributions approximate a uniform?)

This method generates sets of uniform random variables on the interval [0,1]. If the user specifies lower and upper bounds for a variable, the [0,1] samples are mapped to the [lower, upper] interval.

The user must first choose the sequence type:

- `halton` or
- `hammersley`

Then three keywords are used to define the sequence and how it is sampled:

- `prime_base`
- `sequence_start`
- `sequence_leap`

Each of these has defaults, so specification is optional.

Theory

The quasi-Monte Carlo sequences of Halton and Hammersley are deterministic sequences determined by a set of prime bases. Generally, we recommend that the user leave the default setting for the bases, which are the lowest primes. Thus, if one wants to generate a sample set for 3 random variables, the default bases used are 2, 3, and 5 in the Halton sequence. To give an example of how these sequences look, the Halton sequence in base 2 starts with points 0.5, 0.25, 0.75, 0.125, 0.625, etc. The first few points in a Halton base 3 sequence are 0.33333, 0.66667, 0.11111, 0.44444, 0.77777, etc. Notice that the Halton sequence tends to alternate back and forth, generating a point closer to zero then a point closer to one. An individual sequence is based on a radix inverse function defined on a prime base. The prime base determines how quickly the [0,1] interval is filled in. Generally, the lowest primes are recommended.

The Hammersley sequence is the same as the Halton sequence, except the values for the first random variable are equal to 1/N, where N is the number of samples. Thus, if one wants to generate a sample set of 100 samples for 3 random variables, the first random variable has values 1/100, 2/100, 3/100, etc. and the second and third variables are generated according to a Halton sequence with bases 2 and 3, respectively.

For more information about these sequences, see [[43] "Halton, 1960", [44] "Halton and Smith, 1964", and [1] "Kocis and Whiten, 1997"].

See Also

These keywords may also be of interest:

- dace
- fsu_cvt
- psuade_moat

halton

- Keywords Area
- method
- fsu_quasi_mc
- halton

Generate samples from a Halton sequence

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The quasi-Monte Carlo sequences of Halton are deterministic sequences determined by a set of prime bases.

Generally, we recommend that the user leave the default setting for the bases, which are the lowest primes. Thus, if one wants to generate a sample set for 3 random variables, the default bases used are 2, 3, and 5 in the Halton sequence. To give an example of how these sequences look, the Halton sequence in base 2 starts with points 0.5, 0.25, 0.75, 0.125, 0.625, etc. The first few points in a Halton base 3 sequence are 0.33333, 0.66667, 0.11111, 0.44444, 0.77777, etc. Notice that the Halton sequence tends to alternate back and forth, generating a point closer to zero then a point closer to one. An individual sequence is based on a radix inverse function defined on a prime base. The prime base determines how quickly the [0,1] interval is filled in.

Theory

For more information about these sequences, see [[43] "Halton, 1960", [44] "Halton and Smith, 1964", and [1] "Kocis and Whiten, 1997"].

hammersley

- Keywords Area
- method
- fsu_quasi_mc
- hammersley

Use Hammersley sequences

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

See the parent page

latinize

- Keywords Area
- method
- fsu_quasi_mc
- latinize

Adjust samples to improve the discrepancy of the marginal distributions

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The `latinize` control takes the samples and "latinizes" them, meaning that each original sample is moved so that it falls into one strata or bin in each dimension as in Latin Hypercube sampling. The default setting is NOT to latinize. However, one may be interested in doing this in situations where one wants better discrepancy of the 1-dimensional projections (the marginal distributions).

quality_metrics

- Keywords Area
- method
- fsu_quasi_mc
- quality_metrics

Calculate metrics to assess the quality of the samples

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

`quality_metrics` calculates four quality metrics relating to the volumetric spacing of the samples. The four quality metrics measure different aspects relating to the uniformity of point samples in hypercubes. Desirable properties of such point samples are:

- are the points equally spaced
- do the points cover the region
- and are they isotropically distributed
- with no directional bias in the spacing

The four quality metrics we report are:

- h: the point distribution norm, which is a measure of uniformity of the point distribution
- chi: a regularity measure, and provides a measure of local uniformity of a set of points
- tau: the second moment trace measure
- d: the second moment determinant measure

All of these values are scaled so that smaller is better (the smaller the metric, the better the uniformity of the point distribution).

Examples

Complete explanation of these measures can be found in [[38] "Gunzburger and Burkardt, 2004."].

variance_based_decomp

- Keywords Area
- method
- fsu_quasi_mc
- variance_based_decomp

Global sensitivity analyis using variance-based decomposition

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		drop_tolerance	Drop values lower than this tolerance

Description

Dakota can calculate sensitivity indices through variance based decomposition using the keyword `variance_based_decomp`. These indicate how important the uncertainty in each input variable is in contributing to the output variance.

`variance_based_decomp` is extremely computationally intensive since replicated sets of sample values are evaluated. If the user specified a number of samples, N, and a number of nondeterministic variables, M, variance-based decomposition requires the evaluation of N∗(M+2) samples. **Note that specifying this keyword will increase the number of function evaluations above the number requested with the `samples` keyword** .

Because of the computational cost, `variance_based_decomp` is turned off as a default.

To obtain sensitivity indices that are reasonably accurate, we recommend that N, the number of samples, be at least one hundred and preferably several hundred or thousands.

Theory

In this context, we take sensitivity analysis to be global, not local as when calculating derivatives of output variables with respect to input variables. Our definition is similar to that of [[74] "Saltelli et al., 2004"]: "The study of how uncertainty in the output of a model can be apportioned to different sources of uncertainty in the model input."

Variance based decomposition is a way of using sets of samples to understand how the variance of the output behaves, with respect to each input variable. A larger value of the sensitivity index, S_i, means that the uncertainty in the input variable i has a larger effect on the variance of the output. More details on the calculations and interpretation of the sensitivity indices can be found in [[74] "Saltelli et al., 2004"] and [[88] "Weirs et al., 2010"].

drop_tolerance

- Keywords Area
- method
- fsu_quasi_mc
- variance_based_decomp
- drop_tolerance

Drop values lower than this tolerance

Specification

Alias: none

Argument(s): REAL

Description

Drop values lower than this tolerance. For PCE which outputs main, interaction, and total effects by default, the univariate_effects option has been added to allow suppression of the interaction effects since the output volume of these results can be prohibitive for high dimensional problems. Similar to suppression of these interactions is the covariance control, which can be selected to be diagonal_covariance or full_covariance, with the former supporting suppression of the off-diagonal covariance terms (to save compute and memory resources and reduce output volume).

samples

- Keywords Area
- method
- fsu_quasi_mc
- samples

Specify the number of samples

Specification

Alias: none
Argument(s): INTEGER

Description

Specify the number of samples

fixed_sequence

- Keywords Area
- method
- fsu_quasi_mc
- fixed_sequence

Reuse the same sequence and samples for multiple sampling sets

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): none

Description

The `fixed_sequence` control is similar to `fixed_seed` for other sampling methods. If `fixed_sequence` is specified, the user will get the same sequence (meaning the same set of samples) for subsequent calls of the QMC sampling method (for example, this might be used in a surrogate based optimization method or a parameter study where one wants to fix the uncertain variables).

sequence_start

- Keywords Area
- method
- fsu_quasi_mc
- sequence_start

Choose where to start sampling the sequence

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGERLIST

Description

`sequence_start` determines where in the sequence the samples will start.

The default `sequence_start` is a vector with 0 for each variable, specifying that each sequence start with the first term.

Examples

For example, for the Halton sequence in base 2, if the user specifies sequence_start = 2, the sequence would not include 0.5 and 0.25, but instead would start at 0.75.

sequence_leap

- Keywords Area
- method
- fsu_quasi_mc
- sequence_leap

Specify how often the sequence is sampled

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGERLIST

Description

`sequence_leap` controls the "leaping" of terms in the sequence. The default is 1 for each variable, meaning that each term in the sequence be returned.

Examples

If the user specifies a sequence_leap of 2 for a variable, the points returned would be every other term from the QMC sequence.

Theory

The advantage to using a leap value greater than one is mainly for high-dimensional sets of random deviates. In this case, setting a leap value to the next prime number larger than the largest prime base can help maintain uniformity when generating sample sets for high dimensions. For more information about the efficacy of leaped Halton sequences, see [[72] "Robinson and Atcitty, 1999"].

prime_base

- Keywords Area
- method
- fsu_quasi_mc
- prime_base

The prime numbers used to generate the sequence

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGERLIST

Description

It is recommended that the user not specify this and use the default values.

- For the Halton sequence, the default bases are primes in increasing order, starting with 2, 3, 5, etc.
- For the Hammersley sequence, the user specifies (s-1) primes if one is generating an s-dimensional set of random variables.

model_pointer

- Keywords Area
- method
- fsu_quasi_mc
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.70 vector_parameter_study

- Keywords Area
- method
- vector_parameter_study

Parameter study - samples points along a vector

Topics

This keyword is related to the topics:

- parameter_studies

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	final_point	Define the vector by initial and final points
			step_vector	Define the vector with the inital point plus direction and magnitude of each step
	Required		num_steps	Select the number of sampling steps along the vector
	Optional		model_pointer	Specify which model block will be used by a method

Description

Dakota's vector parameter study computes response data sets at selected intervals along a vector in parameter space. It is often used for single-coordinate parameter studies (to study the effect of a single variable on a response set), but it can be used more generally for multiple coordinate vector studies (to investigate the response variations along some n-dimensional vector such as an optimizer search direction).

Group 1 is used to define the vector along which the parameters are varied. Both cases also rely on the variables specification of an initial value, through:

- the initial_point keyword
- the initial_state keyword
- relying on the default initial value, based on the rest of the variables specification

From the initial value, the vector can be defined using one of the two keyword choices.

Once the vector is defined, the samples are then fully specifed by num_steps.

Examples

The following example is a good comparison to the examples on multidim_parameter_study and centered_parameter-_study.

```
# tested on Dakota 6.0 on 140501
environment
  tabular_graphics_data
    tabular_graphics_file = 'rosen_vector.dat'

method
  vector_parameter_study
    num_steps = 10
    final_point =    2.0      2.0
model
  single

variables
  continuous_design = 2
```

```
    initial_point =   -2.0     -2.0
    descriptors =     'x1'     "x2"

interface
  analysis_driver = 'rosenbrock'
    fork

responses
  response_functions = 1
  no_gradients
  no_hessians
```

See Also

These keywords may also be of interest:

- centered_parameter_study
- multidim_parameter_study
- list_parameter_study

final_point

- Keywords Area
- method
- vector_parameter_study
- final_point

Define the vector by initial and final points

Specification

Alias: none
Argument(s): REALLIST

Description

List the final values for each variable. The vector's direction and magnitude are determined by the initial value from the variables specification, and the `final_point`.

The actual points are determined based on the this vector and the number of steps. of points chosen in `num_points`.

step_vector

- Keywords Area
- method
- vector_parameter_study
- step_vector

Define the vector with the inital point plus direction and magnitude of each step

Specification

Alias: none
Argument(s): REALLIST

Description

The step vector defines the change in value for every variable, for a single step. Step $n + 1$ is step n plus the step vector.

The points are determined by this step vector, plus the number of steps.

num_steps

- Keywords Area
- method
- vector_parameter_study
- num_steps

Select the number of sampling steps along the vector

Specification

Alias: none
Argument(s): INTEGER

Description

`num_steps` defines the number of steps that are taken in the direction of the vector. The magnitude of each step is determined in conjuction with the rest of the method specification.

This study performs function evaluations at both ends, making the total number of evaluations equal to `num_steps+1`.

The study has stringent requirements on performing appropriate steps with any discrete range and discrete set variables. A `num_steps` specification must result in discrete range and set index steps that are integers: no remainder is currently permitted in the integer step calculation and no rounding to integer steps will occur.

model_pointer

- Keywords Area
- method
- vector_parameter_study
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.71 list_parameter_study

- Keywords Area
- method
- list_parameter_study

Parameter study specified by a list of variable values

Topics

This keyword is related to the topics:

- parameter_studies

Specification

Alias: none
Argument(s): none

	Required/-Optional	**Description of Group**	**Dakota Keyword**	**Dakota Keyword Description**
	Required*(Choose One)*	**Group 1**	list_of_points	Specify the list of points that the `list_parameter_study` will run
			import_points_file	Specify file containing points (variable and response values)
	Optional		model_pointer	Specify which model block will be used by a method

Description

Dakota's list parameter study allows for evaluations at user selected points of interest.

The number of real values in the `list_of_points` specification or file referenced by `import_points_file` must be a multiple of the total number of variables (including continuous and discrete types) contained in the variables specification.

Notes

- This parameter study simply performs simulations for the first parameter set (the first `n` entries in the list), followed by the next parameter set (the next `n` entries), and so on, until the list of points has been exhausted.
- Since the initial values from the variables specification will not be used, they need not be specified.
- For discrete set types, the actual values should be specified, not the set indices, although the values will be validated for membership within the set value specifications.

Examples

This shows the method and variables block of a Dakota input file that runs a list_parameter_study.

```
method
  list_parameter_study
    list_of_points =
                 3.1e6    0.0029    0.31
                 3.2e6    0.0028    0.32
                 3.3e6    0.0027    0.34
                 3.3e6    0.0026    0.36

variables
  continuous_design = 3
    descriptors = 'E'    'MASS'     'DENSITY'
```

Note that because of the way Dakota treats whitespace, the above example is equivalent to:

```
method
  list_parameter_study
    list_of_points =
3.1e6    0.0029    0.31  3.2e6    0.0028
0.32    3.3e6    0.0027
0.34  3.3e6    0.0026    0.36

variables
  continuous_design = 3
    descriptors = 'E'  'MASS'   'DENSITY'
```

Although the first example is much more readable.

And here's a full input file:

```
# tested on Dakota 6.0 on 140501
environment
  tabular_graphics_data
    tabular_graphics_file 'List_param_study.dat'

method
  list_parameter_study
    list_of_points =     0.1    0.1
                         0.2    0.1
                         0.3    0.0
                         0.3    1.0
```

```
model
  single

variables
  active design
  continuous_design = 2
    descriptors    'x1' 'x2'
  continuous_state = 1
    descriptors =   'constant1'
    initial_state = 100

interface
  analysis_drivers 'text_book'
    fork
  asynchronous
    evaluation_concurrency 2

responses
  response_functions = 1
  no_gradients
  no_hessians
```

This example illustrates the list_parameter_study.

- The function evaluations are independent, so any level of `evaluation_concurrency` can be used
- Default behavior for parameter studies is to iterate on all variables. However, because `active design` is specified, this study will only iterate on the `continuous_design` variables.

See Also

These keywords may also be of interest:

- centered_parameter_study
- multidim_parameter_study
- vector_parameter_study

list_of_points

- Keywords Area
- method
- list_parameter_study
- list_of_points

Specify the list of points that the `list_parameter_study` will run

Specification

Alias: none
Argument(s): REALLIST

Description

The list of points is specified in the freeform format. Dakota will check to see that the number of values in the list is an integer multiple of the number of variables.

import_points_file

- Keywords Area
- method
- list_parameter_study
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: none

Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- method
- list_parameter_study
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- method
- list_parameter_study
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

model_pointer

- Keywords Area
- method
- list_parameter_study
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.72 centered_parameter_study

- Keywords Area
- method
- centered_parameter_study

Parameter study - points moving out from a center point

Topics

This keyword is related to the topics:

- parameter_studies

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		step_vector	Define the size of steps to be taken in each dimension
	Required		steps_per_variable	Number of steps to take in each dimension
	Optional		model_pointer	Specify which model block will be used by a method

Description

Dakota's centered parameter study computes response data sets along multiple coordinate-based vectors, one per parameter, centered about the initial values from the variables specification. This is useful for investigation of function contours with respect to each parameter individually in the vicinity of a specific point (e.g., post-optimality analysis for verification of a minimum), thereby avoiding the cost associated with a multidimensional grid.

In contrast to the `vector_parameter_study`, where steps are taken along a vector, the `centered_parameter_study` takes steps along each orthogonal dimension. Each dimension is treated independently.

The number of steps are taken in each direction, so that the total number of points in the parameter study is $1 + 2\sum n$.

Examples

The following example is a good comparison to the examples on multidim_parameter_study and vector_parameter_study.

```
# tested on Dakota 6.0 on 140501
environment
  tabular_graphics_data
    tabular_graphics_file = 'rosen_centered.dat'

method
  centered_parameter_study
    steps_per_variable = 5 4
    step_vector = 0.4 0.5

model
  single

variables
```

```
  continuous_design = 2
    initial_point =   0        0
    descriptors =     'x1'     "x2"

interface
  analysis_driver = 'rosenbrock'
    fork

responses
  response_functions = 1
  no_gradients
  no_hessians
```

See Also

These keywords may also be of interest:

- multidim_parameter_study
- list_parameter_study
- vector_parameter_study

step_vector

- Keywords Area
- method
- centered_parameter_study
- step_vector

Define the size of steps to be taken in each dimension

Specification

Alias: none
Argument(s): REALLIST

Description

In contrast to the `step_vector` in a `vector_parameter_study`, which defines a vector in n-dimensions that applies to each individual step, this keyword defines the individual step size in each dimension, treated separately.

steps_per_variable

- Keywords Area
- method
- centered_parameter_study
- steps_per_variable

Number of steps to take in each dimension

Specification

Alias: deltas_per_variable
Argument(s): INTEGERLIST

Description

Because the steps in each dimesion are taken independently, the number of steps can be specified for each.

Steps are taken in the plus and minus directions, and are defined in either actual values (continuous and discrete range) or index offsets (discrete set).

model_pointer

- Keywords Area
- method
- centered_parameter_study
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.73 multidim_parameter_study

- Keywords Area
- method
- multidim_parameter_study

Parameter study - make a full factorial grid of study points

Topics

This keyword is related to the topics:

- parameter_studies

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		partitions	Specify the grid of points in the study
	Optional		model_pointer	Specify which model block will be used by a method

Description

Dakota's multidimensional parameter study computes response data sets for an n-dimensional grid of points. Each continuous and discrete range variable is partitioned into equally spaced intervals between its upper and lower bounds, each discrete set variable is partitioned into equally spaced index intervals. The partition boundaries in n-dimensional space define a grid of points, and every point is evaluated.

Since the initial values from the variables specification will not be used, they need not be specified.

Examples

This example is taken from the Users Manual and is a good comparison to the examples on centered_parameter_-study and vector_parameter_study.

```
# tested on Dakota 6.0 on 140501
environment
  tabular_graphics_data
    tabular_graphics_file = 'rosen_multidim.dat'

method
  multidim_parameter_study
    partitions = 10 8

model
  single

variables
  continuous_design = 2
    lower_bounds     -2.0     -2.0
    upper_bounds      2.0      2.0
    descriptors       'x1'     "x2"

interface
  analysis_driver = 'rosenbrock'
    fork

responses
  response_functions = 1
  no_gradients
  no_hessians
```

This example illustrates the full factorial combinations of parameter values created by the multidim_parameter-_study. With 10 and 8 partitions, there are actually 11 and 9 values for each variable. This means that $11 \times 9 = 99$ function evaluations will be required.

See Also

These keywords may also be of interest:

- centered_parameter_study
- list_parameter_study
- vector_parameter_study

partitions

- Keywords Area
- method
- multidim_parameter_study
- partitions

Specify the grid of points in the study

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `partitions` list specifies the number of partitions for each variable. The number of grid points in each dimension is the number of partitions + 1.

The number of entries in the partitions list must be equal to the total number of variables contained in the variables specification.

Remainders within the integer division of the step calculations are not permitted for discrete range or set types and therefore no integer rounding occurs, so the partitions specification must be carefully selected in the presence of these types.

model_pointer

- Keywords Area
- method
- multidim_parameter_study
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.2.74 richardson_extrap

- Keywords Area
- method
- richardson_extrap

Estimate order of convergence of a response as model fidelity increases

Specification

Alias: none
Argument(s): none

	Required/-Optional	**Description of Group**	**Dakota Keyword**	**Dakota Keyword Description**
	Required*(Choose One)*	**Group 1**	estimate_order	Compute the best estimate of the convergence order from three points
			converge_order	Refine until the estimated covergence order converges
			converge_qoi	Refine until the response converges
	Optional		refinement_rate	Rate at which the state variables are refined
	Optional		model_pointer	Specify which model block will be used by a method

Description

Solution verification procedures estimate the order of convergence of the simulation response data during the course of a refinement study. This branch of methods is new and currently only contains one algorithm: Richardson extrapolation.

Refinement of the model

The model fidelity must be parameterized by one or more continuous state variable(s).

The refinement path is determined from the `initial_state` of the `continuous_state` variables specification in combination with the `refinement_rate`, where each of the state variables is treated as an independent refinement factor and each of the initial state values is repeatedly divided by the refinement rate value to define new discretization states.

Results

Three algorithm options are currently provided:

1. `estimate_order`
2. `converge_order`
3. `converge_qoi`

Stopping Criteria

The method employs the `max_iterations` and `convergence_tolerance` method independent controls as stopping criteria.

Theory

In each of these cases, convergence order for a response quantity of interest (QoI) is estimated from

$$p = ln\left(\frac{QoI_3 - QoI_2}{QoI_2 - QoI_1}\right)/ln(r)$$

where r is the uniform refinement rate specified by `refinement_rate`.

estimate_order

- Keywords Area
- method
- richardson_extrap
- estimate_order

Compute the best estimate of the convergence order from three points

Specification

Alias: none

Argument(s): none

Description

The `estimate_order` option is the simplest option. For each of the refinement factors, it evaluates three points along the refinement path and uses these results to perform an estimate of the convergence order for each response function.

converge_order

- Keywords Area
- method
- richardson_extrap
- converge_order

Refine until the estimated covergence order converges

Specification

Alias: none
Argument(s): none

Description

The `converge_order` option is initialized using the `estimate_order` aproach, and additional refinements are performed along the refinement path until the convergence order estimates converge (two-norm of the change in response orders is less than the convergence tolerance).

converge_qoi

- Keywords Area
- method
- richardson_extrap
- converge_qoi

Refine until the response converges

Specification

Alias: none
Argument(s): none

Description

The `converge_qoi` option is similar to the `converge_order` option, except that the convergence criterion is that the two-norm of the response discretization errors (computed from extrapolation) must be less than the convergence tolerance.

refinement_rate

- Keywords Area
- method
- richardson_extrap
- refinement_rate

Rate at which the state variables are refined

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): REAL

Description

Described on parent page

model_pointer

- Keywords Area
- method
- richardson_extrap
- model_pointer

Specify which model block will be used by a method

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `model_pointer` is used to specify which model block will be used to perform function evaluations for the Dakota method. The string value is cross-referenced with id_model keyword in the model block.

See block_pointer for details about pointers.

5.3 model

- Keywords Area
- model

Specifies how variables are mapped into a set of responses

Topics

This keyword is related to the topics:

- block
- not_yet_reviewed

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		id_model	Give the model block an identifying name, in case of multiple model blocks
	Optional		variables_pointer	Specify which variables block will be included with this model block
	Optional		responses_pointer	Specify which reponses block will be used by this model block
	Optional		hierarchical_-tagging	Enables hierarchical evaluation tagging
	Required*(Choose One)*	**Group 1**	single	A model with one of each block: variable, interface, and response
			surrogate	An empirical model that is created from data or the results of a submodel
			nested	A model whose responses are computed through the use of a sub-iterator

Description

A model is comprised of a mapping from variables, through an interface, to responses.

Model Group 1 The type of model can be:

1. `single`

2. `nested`

3. `surrogate`

The input file must specify one of these types. If the type is not specified, Dakota will assume a single model.

Block Pointers and ID

Each of these model types supports `variables_pointer` and `responses_pointer` strings for identifying the variables and responses specifications used in constructing the model by cross-referencing with `id_variables` and `id_responses` strings from particular variables and responses keyword specifications.

These pointers are valid for each model type since each model contains a set of variables that is mapped into a set of responses – only the specifics of the mapping differ.

Additional pointers are used for each model type for constructing the components of the variable to response mapping. As an environment specification identifies a top-level method and a method specification identifies a model, a model specification identifies variables, responses, and (for some types) interface specifications. This top-down flow specifies all of the object interrelationships.

Examples

The next example displays a surrogate model specification which selects a quadratic polynomial from among the global approximation methods. It uses a pointer to a design of experiments method for generating the data needed for building the global approximation, reuses any old data available for the current approximation region, and employs the first-order multiplicative approach to correcting the approximation each time correction is requested.

```
model,
  id_model = 'M1'
  variables_pointer = 'V1'
  responses_pointer = 'R1'
  surrogate
    global
      polynomial quadratic
      dace_method_pointer = 'DACE'
      reuse_samples region
      correction multiplicative first_order
```

This example demonstrates the use of identifiers and pointers. It provides the optional model independent specifications for model identifier, variables pointer, and responses pointer as well as model dependent specifications for global surrogates (see global).

Finally, an advanced nested model example would be

```
model
    id_model = 'M1'
    variables_pointer = 'V1'
    responses_pointer = 'R1'
    nested
     optional_interface_pointer = 'OI1'
      optional_interface_responses_pointer = 'OIR1'
     sub_method_pointer = 'SM1'
      primary_variable_mapping  = '' '' 'X'   'Y'
      secondary_variable_mapping = '' '' 'mean' 'mean'
      primary_response_mapping  = 1. 0. 0. 0. 0. 0. 0. 0. 0.
      secondary_response_mapping = 0. 0. 0. 1. 3. 0. 0. 0. 0.
                      0. 0. 0. 0. 0. 0. 1. 3. 0.
```

This example also supplies model independent controls for model identifier, variables pointer, and responses pointer and supplies model dependent controls for specifying details of the nested mapping.

5.3.1 id_model

- Keywords Area
- model
- id_model

Give the model block an identifying name, in case of multiple model blocks

Topics

This keyword is related to the topics:

- block_identifier

Specification

Alias: none
Argument(s): STRING

Description

The model identifier string is supplied with `id_model` and is used to provide a unique identifier string for use within method specifications (refer to any of the keywords: `model_pointer` in under one of the methods in the method block, for example: model_pointer)

This is used to determine which model the method will run.

See Also

These keywords may also be of interest:

- model_pointer

5.3.2 variables_pointer

- Keywords Area
- model
- variables_pointer

Specify which variables block will be included with this model block

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The `variables_pointer` is used to specify which variables block will be used by the model, by cross-referencing with id_variables keyword in the variables block.

See block_pointer for details about pointers.

5.3.3 responses_pointer

- Keywords Area
- model
- responses_pointer

Specify which reponses block will be used by this model block

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

The `responses_pointer` is used to specify which responses block will be used by the model, by cross-referencing with id_responses keyword in the responses block.

See block_pointer for details about pointers.

5.3.4 hierarchical_tagging

- Keywords Area
- model
- hierarchical_tagging

Enables hierarchical evaluation tagging

Specification

Alias: none

Argument(s): none

Description

The hierarchical tagging option is useful for studies involving multiple models with a nested or hierarchical relationship. For example a nested model has a sub-method, which itself likely operates on a sub-model, or a hierarchical approximation involves coordination of low and high fidelity models. Specifying `hierarchical_tagging` will yield function evaluation identifiers ("tags") composed of the evaluation IDs of the models involved, e.g., outermodel.innermodel.interfaceid = 4.9.2. This communicates the outer contexts to the analysis driver when performing a function evaluation.

Examples

test/dakota_uq_timeseries_ivp_optinterf.in test/dakota_uq_timeseries_sop_optinterf.in

See Also

These keywords may also be of interest:

- file_tag model-nested

5.3.5 single

- Keywords Area
- model
- single

A model with one of each block: variable, interface, and response

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		interface_pointer	Interface block pointer for the single model type

Description

The single model is the simplest model type. It uses a single interface instance to map variables into responses. There is no recursion in this case.

The optional interface_pointer specification identifies the interface block by cross-referencing with the `id_interface` string input from a particular interface keyword specification. This is only necessary when the input file has multiple interface blocks, and you wish to explicitly point to the desired block. The same logic follows for responses and variables blocks and pointers.

Examples

The example shows a minimal specification for a single model, which is the default model when no models are specified by the user.

```
model
  single
```

This example does not provide any pointer strings and therefore relies on the default behavior of constructing the model with the last variables, interface, and responses specifications parsed.

See Also

These keywords may also be of interest:

- surrogate
- nested

interface_pointer

- Keywords Area
- model
- single
- interface_pointer

Interface block pointer for the single model type

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

In the `single` model case, a single interface is used to map the variables into responses. The optional `interface_pointer` specification identifies this interface by cross-referencing with the `id_interface` string input from a particular interface keyword specification.

See block_pointer for details about pointers.

5.3.6 surrogate

- Keywords Area
- model
- surrogate

An empirical model that is created from data or the results of a submodel

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		id_surrogates	Identifies the subset of the response functions by number that are to be approximated (the default is all functions).
	Required*(Choose One)*	**Group 1**	global	Select a surrogate model with global support
			multipoint	Construct a surrogate from multiple existing training points
			local	Build a locally accurate surrogate from data at a single point
			hierarchical	Hierarchical approximations use corrected results from a low fidelity model as an approximation to the results of a high fidelity "truth" model.

Description

Surrogate models are inexpensive approximate models that are intended to capture the salient features of an expensive high-fidelity model. They can be used to explore the variations in response quantities over regions of the parameter space, or they can serve as inexpensive stand-ins for optimization or uncertainty quantification studies (see, for example, the surrogate-based optimization methods, surrogate_based_global and surrogate_based_local).

Surrogate models supported in Dakota are categorized as Data Fitting or Hierarchical, as shown below. Each of these surrogate types provides an approximate representation of a "truth" model which is used to perform the parameter to response mappings. This approximation is built and updated using results from the truth model, called the "training data".

- Data fits:

  Data fitting methods involve construction of an approximation or surrogate model using data (response values, gradients, and Hessians) generated from the original truth model. Data fit methods can be further

categorized as local, multipoint, and global approximation techniques, based on the number of points used in generating the data fit.

1. Local: built from response data from a single point in parameter space
   - Taylor series expansion: taylor_series
     Training data consists of a single point, plus gradient and Hessian information.
2. Multipoint: built from two or more points in parameter space, often involving the current and previous iterates of a minimization algorithm.
   - TANA-3: tana
     Training Data comes from a few previously evaluated points
3. Global full space response surface methods:
   - Polynomial regression: polynomial
   - Gaussian process (Kriging): gaussian_process
   - Artifical neutral network: neural_network
   - MARS: mars
   - Radial Basis Functions: radial_basis
   - Orthogonal polynomials (only supported in PCE/SC for now): polynomial_chaos and stoch_-collocation

   Training data is generated using either a design of experiments method applied to the truth model (specified by dace_method_pointer), or from saved data (specified by reuse_points) in a restart database, or an import file.

- Multifidelity/hierarchical:

  Multifidelity modeling involves the use of a low-fidelity physics-based model as a surrogate for the original high-fidelity model. The low-fidelity model typically involves a coarser mesh, looser convergence tolerances, reduced element order, or omitted physics.

  See hierarchical.

The global and hierarchal surrogates have a correction feature in order to improve the local accuracy of the surrogate models. The correction factors force the surrogate models to match the true function values and possibly true function derivatives at the center point of each trust region. Details can be found on global correction or hierarchical correction.

Theory

Surrogate models are used extensively in the surrogate-based optimization and least squares methods, in which the goals are to reduce expense by minimizing the number of truth function evaluations and to smooth out noisy data with a global data fit. However, the use of surrogate models is not restricted to optimization techniques; uncertainty quantification and optimization under uncertainty methods are other primary users.

Data Fit Surrogate Models

A surrogate of the {*data* fit} type is a non-physics-based approximation typically involving interpolation or regression of a set of data generated from the original model. Data fit surrogates can be further characterized by the number of data points used in the fit, where a local approximation (e.g., first or second-order Taylor series) uses data from a single point, a multipoint approximation (e.g., two-point exponential approximations (TPEA) or two-point adaptive nonlinearity approximations (TANA)) uses a small number of data points often drawn from the previous iterates of a particular algorithm, and a global approximation (e.g., polynomial response surfaces, kriging/gaussian_process, neural networks, radial basis functions, splines) uses a set of data points distributed over the domain of interest, often generated using a design of computer experiments.

Dakota contains several types of surface fitting methods that can be used with optimization and uncertainty quantification methods and strategies such as surrogate-based optimization and optimization under uncertainty. These are: polynomial models (linear, quadratic, and cubic), first-order Taylor series expansion, kriging spatial interpolation, artificial neural networks, multivariate adaptive regression splines, radial basis functions, and moving least squares. With the exception of Taylor series methods, all of the above methods listed in the previous sentence are accessed in Dakota through the Surfpack library. All of these surface fitting methods can be applied to problems having an arbitrary number of design parameters. However, surface fitting methods usually are practical only for problems where there are a small number of parameters (e.g., a maximum of somewhere in the range of 30-50 design parameters). The mathematical models created by surface fitting methods have a variety of names in the engineering community. These include surrogate models, meta-models, approximation models, and response surfaces. For this manual, the terms surface fit model and surrogate model are used.

The data fitting methods in Dakota include software developed by Sandia researchers and by various researchers in the academic community.

Multifidelity Surrogate Models

A second type of surrogate is the {*model* hierarchy} type (also called multifidelity, variable fidelity, variable complexity, etc.). In this case, a model that is still physics-based but is of lower fidelity (e.g., coarser discretization, reduced element order, looser convergence tolerances, omitted physics) is used as the surrogate in place of the high-fidelity model. For example, an inviscid, incompressible Euler CFD model on a coarse discretization could be used as a low-fidelity surrogate for a high-fidelity Navier-Stokes model on a fine discretization.

Surrogate Model Selection

This section offers some guidance on choosing from among the available surrogate model types.

- For Surrogate Based Local Optimization, using the surrogate_based_local method with a trust region:

  using the keywords:

  1. surrogate local taylor_series or
  2. surrogate multipoint tana

  will probably work best.

  If for some reason you wish or need to use a global surrogate (not recommended) then the best of these options is likely to be either:

  1. surrogate global gaussian_process surfpack or
  2. surrogate global moving_least_squares.

- For Efficient Global Optimization (EGO), the efficient_global method:

  the default surrogate is: gaussian_process surfpack which is likely to find a more optimal value and/or require fewer true function evaluations than the alternative, gaussian_process dakota. However, the surfpack will likely take more time to build than the `dakota` version. Note that currently the use_derivatives keyword is not recommended for use with EGO based methods.

- For EGO based global interval estimation, the global_interval_est ego method:

  the default gaussian_process surfpack will likely work better than the alternative gaussian_process dakota.

- For Efficient Global Reliability Analysis (EGRA), the global_reliability method:

  the surfpack and `dakota` versions of the gaussian process tend to give similar answers with the `dakota` version tending to use fewer true function evaluations. Since this is based on EGO, it is likely that the default surfpack is more accurate, although this has not been rigorously demonstrated.

- For EGO based Dempster-Shafer Theory of Evidence, i.e. the global_evidence ego method, the default gaussian_process surfpack often use significantly fewer true function evaluations than the alternative gaussian_process dakota.
- When using a global surrogate to extrapolate, any of the surrogates:
  - gaussian_process surfpack
  - polynomial quadratic
  - polynomial cubic

  are recommended.
- When there is over roughly two or three thousand data points and you wish to interpolate (or approximately interpolate) then a Taylor series, Radial Basis Function Network, or Moving Least Squares fit is recommended. The only reason that the gaussian_process surfpack is not recommended is that it can take a considerable amount of time to construct when the number of data points is very large. Use of the third party MARS package included in Dakota is generally discouraged.
- In other situations that call for a global surrogate, the gaussian_process surfpack is generally recommended. The `use_derivatives` keyword will only be useful if accurate and inexpensive derivatives are available. Finite difference derivatives are disqualified on both counts. However, derivatives generated by analytical, automatic differentiation, or continuous adjoint techniques can be appropriate. Currently, first order derivatives, i.e. gradients, are the highest order derivatives that can be used to construct the gaussian_process surfpack model; Hessians will not be used even if they are available.

See Also

These keywords may also be of interest:

- single
- nested

id_surrogates

- Keywords Area
- model
- surrogate
- id_surrogates

Identifies the subset of the response functions by number that are to be approximated (the default is all functions).

Specification

Alias: none
Argument(s): INTEGERLIST

Description

In the `surrogate` model case, the specification first allows a mixture of surrogate and actual response mappings through the use of the optional `id_surrogates` specification. This identifies the subset of the response functions by number that are to be approximated (the default is all functions). The valid response function identifiers range from 1 through the total number of response functions (see response_functions).

global

- Keywords Area
- model
- surrogate
- global

Select a surrogate model with global support

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
			gaussian_process	Use a Gaussian Process surrogate
	Required*(Choose One)*	**Group 1**	mars	Multivariate Adaptive Regression Spline (MARS)
			moving_least_-squares	Moving Least Squares surrogate models
			neural_network	Artificial Neural Network Models
			radial_basis	Radial Basis Functions
			polynomial	Linear, Quadratic, or Cubic Polynomial Models
	Optional*(Choose One)*	**Group 2**	total_points	Specify the number of training points

			minimum_points	Construct surrogate with the minimum amount of points
			recommended_-points	(Default) Construct surrogate with the recommended amount of points
	Optional		dace_method_-pointer	Specify a method to gather training data
	Optional		reuse_points	Specify how to use existing training data
	Optional		import_points_file	Specify file containing points (variable and response values)
	Optional		export_points_file	Specify file to write out the responses from the surrogate
	Optional		use_derivatives	Use derivative data to construct surrogate models
	Optional		correction	Correction approaches for surrogate models
	Optional		metrics	Compute metrics of surrogate quality
	Optional		challenge_points_-file	Specify a datafile of points to compare surrogate responses against

Description

The global surrogate model requires specification of one of the following approximation types:

1. Polynomial
2. Gaussian process (Kriging interpolation)
3. Layered perceptron artificial neural network approximation
4. MARS
5. Moving least squares
6. Radial basis function

All these approximations are implemented in SurfPack[36]. In addition, a second version of Gaussian process is implemented directly in Dakota.

Training Data

Training data can be taken from prior runs, stored in a datafile, or by running a Design of Experiments method. The keywords listed below are used to determine how to collect training data:

- `dace_method_pointer`
- `reuse_points`
- `import_points_file`
- `use_derivatives` The source of training data is determined by the contents of a provided `import_points_file`, whether `reuse_points` and `use_derivatives` are specified, and the contents of the method block specified by `dace_method_pointer`. `use_derivatives` is a special case, the other keywords are discussed below.

The number of training data points used in building a global approximation is determined by specifying one of three point counts:

1. `minimum_points`: minimum "reasonable" amount of training data, based on Dakota's heuristics
2. `recommended_points`: recommended number of training data, based on Dakota's heuristics (this is the default, if none of the keywords is specified)
3. `total_points`: specify the number of training data points. However, if the `total_points` value is less than the default `minimum_points` value, the `minimum_points` value is used.

The sources of training data depend on the number of training points, N_{tp}, the number of points in the import file, N_{if}, and the value of `reuse_points`.

- If there is no import file, all training data come from the DACE method
- If there is an import file, all N_{if} points from the file are used, and the remaining $N_{tp} - N_{if}$ points come from the DACE method
- If there is an import file and `reuse_points` is:
  - `none` - all N_{tp} points from DACE method
  - `region` - only the points within a trust region are taken from the import file, and all remaining points are from the DACE method.
  - `all` - (Default) all N_{if} points from the file are used, and the remaining $N_{tp} - N_{if}$ points come from the DACE method

Surrogate Correction

A `correction` model can be added to the constructed surrogate in order to better match the training data. The specified correction method will be applied to the surrogate, and then the corrected surrogate model is used by the method.

Finally, the quality of the surrogate can be tested using the `metrics` and `challenge_points_file` keywords.

Theory

Global methods, also referred to as response surface methods, involve many points spread over the parameter ranges of interest. These surface fitting methods work in conjunction with the sampling methods and design of experiments methods.

Procedures for Surface Fitting

The surface fitting process consists of three steps:

1. selection of a set of design points
2. evaluation of the true response quantities (e.g., from a user-supplied simulation code) at these design points,
3. using the response data to solve for the unknown coefficients (e.g., polynomial coefficients, neural network weights, kriging correlation factors) in the surface fit model.

In cases where there is more than one response quantity (e.g., an objective function plus one or more constraints), then a separate surface is built for each response quantity. Currently, the surface fit models are built using only 0 th-order information (function values only), although extensions to using higher-order information (gradients and Hessians) are possible.

Each surface fitting method employs a different numerical method for computing its internal coefficients. For example, the polynomial surface uses a least-squares approach that employs a singular value decomposition to compute the polynomial coefficients, whereas the kriging surface uses Maximum Likelihood Estimation to compute its correlation coefficients. More information on the numerical methods used in the surface fitting codes is provided in the Dakota Developers Manual.

See Also

These keywords may also be of interest:

- local
- hierarchical
- multipoint

gaussian_process

- Keywords Area
- model
- surrogate
- global
- gaussian_process

Use a Gaussian Process surrogate

Specification

Alias: kriging

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	dakota	Select the built in Gaussian Process surrogate
			surfpack	Use the Surfpack version of Gaussian Process surrogates

Description

Dakota uses the Gaussian process (GP) surrogate from Surfpack, which is specified using the surfpack keyword.

A second version of GP surrogates was available in prior versions of Dakota.

For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.

dakota

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- dakota

Select the built in Gaussian Process surrogate

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		point_selection	
	Optional		trend	Choose a trend function for a Gaussian process surrogate

Description

A second version of GP surrogates was available in prior versions of Dakota. **For now, both versions are supported but the `dakota` version is deprecated and intended to be removed in a future release.**

Historically these models were drastically different, but in Dakota 5.1, they became quite similar. They now differ in that the Surfpack GP has a richer set of features/options and tends to be more accurate than the Dakota version. Due to how the Surfpack GP handles ill-conditioned correlation matrices (which significantly contributes

to its greater accuracy), the `Surfpack` GP can be a factor of two or three slower than Dakota's. As of Dakota 5.2, the Surfpack implementation is the default in all contexts except Bayesian calibration.

More details on the `gaussian_process` dakota model can be found in[59].

Dakota's GP deals with ill-conditioning in two ways. First, when it encounters a non-invertible correlation matrix it iteratively increases the size of a "nugget," but in such cases the resulting approximation smooths rather than interpolates the data. Second, it has a `point_selection` option (default off) that uses a greedy algorithm to select a well-spaced subset of points prior to the construction of the GP. In this case, the GP will only interpolate the selected subset. Typically, one should not need point selection in trust-region methods because a small number of points are used to develop a surrogate within each trust region. Point selection is most beneficial when constructing with a large number of points, typically more than order one hundred, though this depends on the number of variables and spacing of the sample points.

This differs from the `point_selection` option of the Dakota GP which initially chooses a well-spaced subset of points and finds the correlation parameters that are most likely for that one subset.

point_selection

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- dakota
- point_selection

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none

Argument(s): none

trend

- Keywords Area
- model
- surrogate
- global
- gaussian_process

- dakota
- trend

Choose a trend function for a Gaussian process surrogate

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	constant	Constant term
			linear	Use a linear polynomial
			reduced_quadratic	Quadratic polynomials - main effects only

Description

The only trend functions that are currently supported are polynomials.

The trend function is selected using the `trend` keyword, with options `constant`, `linear`, or `reduced_quadratic`. The `reduced_quadratic` trend function includes the main effects, but not mixed/interaction terms. The Surfpack GP (See surfpack) has the additional option of (a full) `quadratic`.

constant

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- dakota
- trend
- constant

Constant term

Specification

Alias: none

Argument(s): none

Description

See parent page

linear

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- dakota
- trend
- linear

Use a linear polynomial

Specification

Alias: none
Argument(s): none

Description

See parent page

reduced_quadratic

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- dakota
- trend
- reduced_quadratic

Quadratic polynomials - main effects only

Specification

Alias: none
Argument(s): none

Description

In 2+ dimensions, this does not include interactions, a.k.a. mixed terms.

surfpack

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack

Use the Surfpack version of Gaussian Process surrogates

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		trend	Choose a trend function for a Gaussian process surrogate
	Optional		optimization_-method	Change the optimization method used to compute hyperparameters
	Optional		max_trials	Max number of likelihood function evaluations
	Optional(*Choose One*)	**Group 1**	nugget	Specify a nugget to handle ill-conditioning
			find_nugget	Have Surfpack compute a nugget to handle ill-conditioning
	Optional		correlation_lengths	Specify the correlation lengths for the Gaussian process

	Optional	export_model_file	????

Description

This keyword specifies the use of the Gaussian process that is incorporated in our surface fitting library called Surfpack.

Several user options are available:

1. Optimization methods:

   Maximum Likelihood Estimation (MLE) is used to find the optimal values of the hyper-parameters governing the trend and correlation functions. By default the global optimization method DIRECT is used for MLE, but other options for the optimization method are available. See optimization_method.

   The total number of evaluations of the likelihood function can be controlled using the `max_trials` keyword followed by a positive integer. Note that the likelihood function does not require running the "truth" model, and is relatively inexpensive to compute.

2. Trend Function:

   The GP models incorporate a parametric trend function whose purpose is to capture large-scale variations. See trend.

3. Correlation Lengths:

   Correlation lengths are usually optimized by Surfpack, however, the user can specify the lengths manually. See correlation_lengths.

4. Ill-conditioning

   One of the major problems in determining the governing values for a Gaussian process or Kriging model is the fact that the correlation matrix can easily become ill-conditioned when there are too many input points close together. Since the predictions from the Gaussian process model involve inverting the correlation matrix, ill-conditioning can lead to poor predictive capability and should be avoided.

   Note that a sufficiently bad sample design could require correlation lengths to be so short that any interpolatory GP model would become inept at extrapolation and interpolation.

   The `surfpack` model handles ill-conditioning internally by default, but behavior can be modified using

5. Gradient Enhanced Kriging (GEK).

   The `use_derivatives` keyword will cause the Surfpack GP to be constructed from a combination of function value and gradient information (if available).

   See notes in the Theory section.

Theory

Gradient Enhanced Kriging

Incorporating gradient information will only be beneficial if accurate and inexpensive derivative information is available, and the derivatives are not infinite or nearly so. Here "inexpensive" means that the cost of evaluating a function value plus gradient is comparable to the cost of evaluating only the function value, for example gradients computed by analytical, automatic differentiation, or continuous adjoint techniques. It is not cost effective to use derivatives computed by finite differences. In tests, GEK models built from finite difference derivatives were also significantly less accurate than those built from analytical derivatives. Note that GEK's correlation matrix tends to have a significantly worse condition number than Kriging for the same sample design.

This issue was addressed by using a pivoted Cholesky factorization of Kriging's correlation matrix (which is a small sub-matrix within GEK's correlation matrix) to rank points by how much unique information they contain. This reordering is then applied to whole points (the function value at a point immediately followed by gradient information at the same point) in GEK's correlation matrix. A standard non-pivoted Cholesky is then applied to the reordered GEK correlation matrix and a bisection search is used to find the last equation that meets the constraint on the (estimate of) condition number. The cost of performing pivoted Cholesky on Kriging's correlation matrix is usually negligible compared to the cost of the non-pivoted Cholesky factorization of GEK's correlation matrix. In tests, it also resulted in more accurate GEK models than when pivoted Cholesky or whole-point-block pivoted Cholesky was performed on GEK's correlation matrix.

?? This correlation function results in a response surface that is C^{∞}-continuous.

trend

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- trend

Choose a trend function for a Gaussian process surrogate

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword constant	Dakota Keyword Description Constant term
	Required*(Choose One)*	**Group 1**	linear	Use a linear polynomial
			reduced_quadratic	Quadratic polynomials - main effects only
			quadratic	Use a quadratic polynomial

Description

The only trend functions that are currently supported are polynomials.

The trend function is selected using the `trend` keyword, with options `constant`, `linear`, or `reduced_quadratic`. The `reduced_quadratic` trend function includes the main effects, but not mixed/interaction terms. The Surfpack GP (See surfpack) has the additional option of (a full) `quadratic`.

constant

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- trend
- constant

Constant term

Specification

Alias: none
Argument(s): none

Description

See parent page

linear

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- trend
- linear

Use a linear polynomial

Specification

Alias: none
Argument(s): none

Description

See parent page

reduced_quadratic

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- trend
- reduced_quadratic

Quadratic polynomials - main effects only

Specification

Alias: none
Argument(s): none

Description

In 2+ dimensions, this does not include interactions, a.k.a. mixed terms.

quadratic

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- trend
- quadratic

Use a quadratic polynomial

Specification

Alias: none
Argument(s): none

Description

See parent page

optimization_method

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- optimization_method

Change the optimization method used to compute hyperparameters

Specification

Alias: none
Argument(s): STRING

Description

Select the optimization method to compute hyperparameters of the Gaussian Process by specifying one of these arguments:

- `global` (default) - DIRECT method
- `local` - CONMIN method
- `sampling` - generates several random guesses and picks the candidate with greatest likelihood
- `none` - no optimization, pick the center of the feasible region

The `none` option, and the starting location of the `local` optimization, default to the center, in log(correlation length) scale, of the of feasible region.
Surfpack picks a small feasible region of correlation parameters.
Note that we have found the `global` optimization method to be the most robust.

max_trials

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- max_trials

Max number of likelihood function evaluations

Specification

Alias: none
Argument(s): INTEGER

Description

See parent page

nugget

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- nugget

Specify a nugget to handle ill-conditioning

Specification

Alias: none
Argument(s): REAL

Description

By default, the Surfpack GP handles ill-conditioning and does not use a nugget. If the user wishes to specify a nugget, there are two approaches.

- The user can specify the value of a nugget with nugget.
- Have Surfpack find the optimal value of the nugget. This is specified by find_nugget. There are two options for find_nugget.
  - `find_nugget` = 1: assume that the reciprocal condition number of the correlation matrix R, rcondR, is zero and calculate the nugget needed to make the worst case of R not ill-conditioned.
  - `find_nugget` = 2: calculate rcondR, which requires a Cholesky factorization. If rcondR indicates that R is not ill-conditioned, then kriging uses the Cholesky factorization. Otherwise, if rcondR says R is ill conditioned, then kriging will calculate the nugget needed to make the worst case of R not ill conditioned.

`find_nugget` = 1 and 2 are similar, the second option just takes more computation (the initial Cholesky factorization) for larger problems.

find_nugget

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- find_nugget

Have Surfpack compute a nugget to handle ill-conditioning

Specification

Alias: none
Argument(s): INTEGER

Description

By default, the Surfpack GP handles ill-conditioning and does not use a nugget. If the user wishes to specify a nugget, there are two approaches.

- The user can specify the value of a nugget with nugget.
- Have Surfpack find the optimal value of the nugget. This is specified by find_nugget. There are two options for find_nugget.
  - `find_nugget` = 1: assume that the reciprocal condition number of the correlation matrix R, rcondR, is zero and calculate the nugget needed to make the worst case of R not ill-conditioned.
  - `find_nugget` = 2: calculate rcondR, which requires a Cholesky factorization. If rcondR indicates that R is not ill-conditioned, then kriging uses the Cholesky factorization. Otherwise, if rcondR says R is ill conditioned, then kriging will calculate the nugget needed to make the worst case of R not ill conditioned.

  `find_nugget` = 1 and 2 are similar, the second option just takes more computation (the initial Cholesky factorization) for larger problems.

correlation_lengths

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- correlation_lengths

Specify the correlation lengths for the Gaussian process

Specification

Alias: none
Argument(s): REALLIST

Description

Directly specify `correlation_lengths` as a list of N real numbers where N is the number of input dimensions.

export_model_file

- Keywords Area
- model
- surrogate
- global
- gaussian_process
- surfpack
- export_model_file

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): STRING

Description

Not mentioned

mars

- Keywords Area
- model
- surrogate
- global
- mars

Multivariate Adaptive Regression Spline (MARS)

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		max_bases	The maximum number of basis allowed to be used in the approximation.
	Optional		interpolation	The interpolation type. Linear or cubic.
	Optional		export_model_file	????

Description

This surface fitting method uses multivariate adaptive regression splines from the MARS3.5 package[26] developed at Stanford University.

The MARS reference material does not indicate the minimum number of data points that are needed to create a MARS surface model. However, in practice it has been found that at least $n_{c_{quad}}$, and sometimes as many as 2 to 4 times $n_{c_{quad}}$, data points are needed to keep the MARS software from terminating. Provided that sufficient data samples can be obtained, MARS surface models can be useful in SBO and OUU applications, as well as in the prediction of global trends throughout the parameter space.

Theory

The form of the MARS model is based on the following expression:

$$\hat{f}(\mathbf{x}) = \sum_{m=1}^{M} a_m B_m(\mathbf{x})$$

where the a_m are the coefficients of the truncated power basis functions B_m, and M is the number of basis functions. The MARS software partitions the parameter space into subregions, and then applies forward and backward regression methods to create a local surface model in each subregion. The result is that each subregion contains its own basis functions and coefficients, and the subregions are joined together to produce a smooth, C^2-continuous surface model.

MARS is a nonparametric surface fitting method and can represent complex multimodal data trends. The regression component of MARS generates a surface model that is not guaranteed to pass through all of the response data values. Thus, like the quadratic polynomial model, it provides some smoothing of the data.

max_bases

- Keywords Area
- model
- surrogate
- global

- mars
- max_bases

The maximum number of basis allowed to be used in the approximation.

Specification

Alias: none
Argument(s): INTEGER

Description

The maximum number of basis allowed to be used in the approximation.

interpolation

- Keywords Area
- model
- surrogate
- global
- mars
- interpolation

The interpolation type. Linear or cubic.

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	linear	Use a linear polynomial
			cubic	Use a cubic polynomial

Description

The interpolation type. Linear or cubic. Not mentioned

linear

- Keywords Area
- model
- surrogate
- global
- mars
- interpolation
- linear

Use a linear polynomial

Specification

Alias: none
Argument(s): none

Description

See parent page

cubic

- Keywords Area
- model
- surrogate
- global
- mars
- interpolation
- cubic

Use a cubic polynomial

Specification

Alias: none
Argument(s): none

Description

See parent page

export_model_file

- Keywords Area
- model
- surrogate
- global
- mars
- export_model_file

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): STRING

Description

Not mentioned

moving_least_squares

- Keywords Area
- model
- surrogate
- global
- moving_least_squares

Moving Least Squares surrogate models

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Optional	poly_order	The degree of thep total degree polynomial basis.
	Optional	weight_function	The weight function used to weight samples in the weighted least squares formulation.
	Optional	export_model_file	????

Description

Moving least squares is a further generalization of weighted least squares where the weighting is "moved" or recalculated for every new point where a prediction is desired[65].

The implementation of moving least squares is still under development. We have found that it works well in trust region methods where the surrogate model is constructed in a constrained region over a few points. It does not appear to be working as well globally, at least at this point in time.

Theory

Moving Least Squares can be considered a more specialized version of linear regression models. In linear regression, one usually attempts to minimize the sum of the squared residuals, where the residual is defined as the difference between the surrogate model and the true model at a fixed number of points.

In weighted least squares, the residual terms are weighted so the determination of the optimal coefficients governing the polynomial regression function, denoted by $\hat{f}(\mathbf{x})$, are obtained by minimizing the weighted sum of squares at N data points:

$$\sum_{n=1}^{N} w_n(\| \hat{f}(\mathbf{x_n}) - f(\mathbf{x_n}) \|)$$

poly_order

- Keywords Area
- model
- surrogate
- global
- moving_least_squares
- poly_order

The degree of thep total degree polynomial basis.

Specification

Alias: none

Argument(s): INTEGER

Description

The degree of thep total degree polynomial basis.

weight_function

- Keywords Area
- model
- surrogate
- global
- moving_least_squares
- weight_function

The weight function used to weight samples in the weighted least squares formulation.

Specification

Alias: none
Argument(s): INTEGER

Description

The weight function used to weight samples in the weighted least squares formulation.

export_model_file

- Keywords Area
- model
- surrogate
- global
- moving_least_squares
- export_model_file

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): STRING

Description

Not mentioned

neural_network

- Keywords Area
- model
- surrogate
- global
- neural_network

Artificial Neural Network Models

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		nodes	????
	Optional		range	????
	Optional		random_weight	????
	Optional		export_model_file	????

Description

The Artificial Neural Network surrogate in Dakota employs a stochastic layered perceptron artificial neural network based on the direct training approach of Zimmerman[93].

The method is designed to have a lower training cost than traditional neural networks. This is a useful feature for surrogate-based optimization and optimization under uncertainty, where new surrogates are constructed many times during the optimization process (i.e., one surrogate for each response function, and new surrogates for each optimization iteration).

The nueral network is a non parametric surface fitting method. Thus, along with kriging (Gaussian Process) and MARS, it can be used to model data trends that have slope discontinuities as well as multiple maxima and minima. However, unlike kriging, the neural network surrogate is not guaranteed to exactly match the response values of the data points from which it was constructed.

This surrogate can be constructed from fewer than $n_{c_{quad}}$ data points, however, it is a good rule of thumb to use at least $n_{c_{quad}}$ data points when possible.

Theory

The form of the SLP ANN model is

$$\hat{f}(\mathbf{x}) \approx \tanh(\tanh((\mathbf{x}\mathbf{A}_0 + \theta_0)\mathbf{A}_1 + \theta_1))$$

where $\mathbf{x}$ is the current point in n-dimensional parameter space, and the terms $\mathbf{A}_0, \theta_0, \mathbf{A}_1, \theta_1$ are the matrices and vectors that correspond to the neuron weights and offset values in the ANN model. These terms are computed during the ANN training process, and are analogous to the polynomial coefficients in a quadratic surface fit. A singular value decomposition method is used in the numerical methods that are employed to solve for the weights and offsets.

nodes

- Keywords Area
- model
- surrogate
- global
- neural_network
- nodes

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): INTEGER

Description

Not mentioned

range

- Keywords Area
- model
- surrogate
- global
- neural_network
- range

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): REAL

Description

Not mentioned

random_weight

- Keywords Area
- model
- surrogate
- global
- neural_network
- random_weight

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): INTEGER

Description

Not mentioned

export_model_file

- Keywords Area
- model
- surrogate
- global
- neural_network
- export_model_file

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): STRING

Description

Not mentioned

radial_basis

- Keywords Area
- model
- surrogate
- global
- radial_basis

Radial Basis Functions

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		bases	
	Optional		max_pts	????
	Optional		min_partition	????
	Optional		max_subsets	????
	Optional		export_model_file	????

Description

Radial basis functions are functions whose value typically depends on the distance from a center point, called the centroid, **c**.

The surrogate model approximation is then built up as the sum of K weighted radial basis functions:

$$\hat{f}(\mathbf{x}) = \sum_{k=1}^{K} w_k \phi(\| \mathbf{x} - \mathbf{c_k} \|)$$

where the ϕ are the individual radial basis functions.

These functions can be of any form, but often a Gaussian bell-shaped function or splines are used.

Our implementation uses a Gaussian radial basis function. The weights are determined via a linear least squares solution approach. See[68] for more details.

bases

- Keywords Area
- model
- surrogate
- global
- radial_basis
- bases

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): INTEGER

max_pts

- Keywords Area
- model
- surrogate
- global
- radial_basis
- max_pts

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): INTEGER

Description

Not mentioned

min_partition

- Keywords Area
- model
- surrogate
- global
- radial_basis
- min_partition

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): INTEGER

Description

Not mentioned

max_subsets

- Keywords Area
- model
- surrogate
- global
- radial_basis
- max_subsets

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): INTEGER

Description

Not mentioned

export_model_file

- Keywords Area
- model
- surrogate
- global
- radial_basis
- export_model_file

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): STRING

Description

Not mentioned

polynomial

- Keywords Area
- model
- surrogate
- global
- polynomial

Linear, Quadratic, or Cubic Polynomial Models

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	linear	Use a linear polynomial
			quadratic	Use a quadratic polynomial
			cubic	Use a cubic polynomial
	Optional		export_model_file	????

Description

Linear, quadratic, or cubic polynomial surrogates models are available in Dakota. The utility of the polynomial models stems from two sources:

- over a small portion of the parameter space, a low-order polynomial model is often an accurate approximation to the true data trends
- the least-squares procedure provides a surface fit that smooths out noise in the data.

The surrogate-based optimization methods are often successful when using polynomial models, particularly quadratic models. However, a polynomial surface fit may not be the best choice for modeling data trends over the entire parameter space, unless it is known a priori that the true data trends are close to linear, quadratic, or cubic. See[64] for more information on polynomial models.

Theory

The form of the linear polynomial model is

$$\hat{f}(\mathbf{x}) \approx c_0 + \sum_{i=1}^{n} c_i x_i$$

the form of the quadratic polynomial model is:

$$\hat{f}(\mathbf{x}) \approx c_0 + \sum_{i=1}^{n} c_i x_i + \sum_{i=1}^{n} \sum_{j \geq i}^{n} c_{ij} x_i x_j$$

and the form of the cubic polynomial model is:

$$\hat{f}(\mathbf{x}) \approx c_0 + \sum_{i=1}^{n} c_i x_i + \sum_{i=1}^{n} \sum_{j \geq i}^{n} c_{ij} x_i x_j + \sum_{i=1}^{n} \sum_{j \geq i}^{n} \sum_{k \geq j}^{n} c_{ijk} x_i x_j x_k$$

In all of the polynomial models, $\hat{f}(\mathbf{x})$ is the response of the polynomial model; the x_i, x_j, x_k terms are the components of the n-dimensional design parameter values; the c_0 , c_i , c_{ij} , c_{ijk} terms are the polynomial coefficients, and n is the number of design parameters. The number of coefficients, n_c, depends on the order of polynomial model and the number of design parameters. For the linear polynomial:

$$n_{c_{linear}} = n + 1$$

for the quadratic polynomial:

$$n_{c_{quad}} = \frac{(n+1)(n+2)}{2}$$

and for the cubic polynomial:

$$n_{c_{cubic}} = \frac{(n^3 + 6n^2 + 11n + 6)}{6}$$

There must be at least n_c data samples in order to form a fully determined linear system and solve for the polynomial coefficients. In Dakota, a least-squares approach involving a singular value decomposition numerical method is applied to solve the linear system.

linear

- Keywords Area
- model
- surrogate
- global
- polynomial
- linear

Use a linear polynomial

Specification

Alias: none
Argument(s): none

Description

See parent page

quadratic

- Keywords Area
- model
- surrogate
- global
- polynomial
- quadratic

Use a quadratic polynomial

Specification

Alias: none
Argument(s): none

Description

See parent page

cubic

- Keywords Area
- model
- surrogate
- global
- polynomial
- cubic

Use a cubic polynomial

Specification

Alias: none
Argument(s): none

Description

See parent page

export_model_file

- Keywords Area
- model
- surrogate
- global
- polynomial
- export_model_file

????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): STRING

Description

Not mentioned

total_points

- Keywords Area
- model
- surrogate
- global
- total_points

Specify the number of training points

Specification

Alias: none
Argument(s): INTEGER

Description

See parent page.

minimum_points

- Keywords Area
- model
- surrogate
- global
- minimum_points

Construct surrogate with the minimum amount of points

Specification

Alias: none
Argument(s): none

Description

See parent page.

FAQ

What is the heuristic?

recommended_points

- Keywords Area
- model
- surrogate
- global
- recommended_points

(Default) Construct surrogate with the recommended amount of points

Specification

Alias: none
Argument(s): none

Description

See parent page.

FAQ

What is the heuristic?

dace_method_pointer

- Keywords Area
- model
- surrogate
- global
- dace_method_pointer

Specify a method to gather training data

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The number of training points and the sources are specified on global, as well as the number of new training points required.

New training points are gathered by running the "truth" model using the method specified by `dace_method_pointer`. The DACE method will only be invoked if it has new samples to perform, and if new samples are required and no DACE iterator has been provided, an error will result.

The `dace_method_pointer` points to design of experiments method block used to generate truth model data.

Permissible methods include: Monte Carlo (random) sampling, Latin hypercube sampling, orthogonal array sampling, central composite design sampling, and Box-Behnken sampling.

Note that the number of samples specified in the method block may be overwritten, if the requested number of samples is less than minimum_points.

reuse_points

- Keywords Area
- model
- surrogate
- global
- reuse_points

Specify how to use existing training data

Topics

This keyword is related to the topics:

- problem

Specification

Alias: reuse_samples

Argument(s): none

	Required/-Optional **Required***(Choose One)*	Description of Group **Group 1**	Dakota Keyword all	Dakota Keyword Description Option for `reuse_points`
			region	Option for `reuse_points`
			none	Option for `reuse_points`

Description

Dakota's global surrogate methods rely on training data, which can either come from evaluation of a "truth" model, which is gathered by the method specified with dace_method_pointer, OR from a file of existing training data, identified by import_points_file.

If `import_points_file` is specified, `reuse_points` modifies how the file contents are used.

- `all` (default) - use all points in the file
- `region` - use only the points in the trust region
- `none` - ignore the contents of the file, and gather new training data using the specified method

FAQ

What is a trust region?

all

- Keywords Area
- model
- surrogate
- global
- reuse_points
- all

Option for `reuse_points`

Specification

Alias: none
Argument(s): none

Description

This is described on the parent page.

region

- Keywords Area
- model
- surrogate
- global
- reuse_points
- region

Option for `reuse_points`

Specification

Alias: none
Argument(s): none

Description

This is described on the parent page.

none

- Keywords Area
- model
- surrogate
- global
- reuse_points
- none

Option for `reuse_points`

Specification

Alias: none
Argument(s): none

Description

This is described on the parent page.

import_points_file

- Keywords Area
- model
- surrogate
- global
- import_points_file

Specify file containing points (variable and response values)

Specification

Alias: samples_file
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format

			freeform	Specified file is in the freeform format

Description

File of points (variable and response values) that will be used by the method in place of function evaluations.

When used to construct surrogate models or emulators these are often called build points or training data.

Although Dakota parses input files without regard to whitespace, the import_points_file must be in one of two formats:

- `annotated` (default)
- `freeform`

annotated

- Keywords Area
- model
- surrogate
- global
- import_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none

Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.

- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- model
- surrogate
- global
- import_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

export_points_file

- Keywords Area
- model
- surrogate
- global
- export_points_file

Specify file to write out the responses from the surrogate

Specification

Alias: none

Argument(s): STRING

	Required/-Optional **Optional***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** annotated	**Dakota Keyword Description** (Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

File of points (inputs and predicted outputs from the surrogate) evaluated on the surrogate model. Note that the export points contain test points and the emulator predictions at these points.

annotated

- Keywords Area
- model
- surrogate
- global
- export_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- model
- surrogate
- global
- export_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

use_derivatives

- Keywords Area
- model
- surrogate
- global
- use_derivatives

Use derivative data to construct surrogate models

Specification

Alias: none

Argument(s): none

Description

The `use_derivatives` flag specifies that any available derivative information should be used in global approximation builds, for those global surrogate types that support it (currently, polynomial regression and the Surfpack Gaussian process).

However, it's use with Surfpack Gaussian process is not recommended.

correction

- Keywords Area
- model
- surrogate
- global
- correction

Correction approaches for surrogate models

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	zeroth_order	Specify that truth values must be matched.
			first_order	Specify that truth values and gradients must be matched.
			second_order	Specify that truth values, gradients and Hessians must be matched.
	Required*(Choose One)*	**Group 2**	additive	Use an additive correction factor to improve the local accuracy of a surrogate.
			multiplicative	Use a multiplicative correction factor to improve the local accuracy of a surrogatel.
			combined	A multipoint correction for a hierarchical surrogate.

Description

Some of the surrogate model types support the use of correction factors that improve the local accuracy of the surrogate models.

The `correction` specification specifies that the approximation will be corrected to match truth data, either matching truth values in the case of `zeroth_order` matching, matching truth values and gradients in the case of `first_order` matching, or matching truth values, gradients, and Hessians in the case of `second_order` matching. For `additive` and `multiplicative` corrections, the correction is local in that the truth data is matched at a single point, typically the center of the approximation region. The `additive` correction adds a scalar offset (`zeroth_order`), a linear function (`first_order`), or a quadratic function (`second_order`) to the approximation to match the truth data at the point, and the `multiplicative` correction multiplies the approximation by a scalar (`zeroth_order`), a linear function (`first_order`), or a quadratic function (`second_order`) to match the truth data at the point. The `additive first_order` case is due to [[58] "-Lewis and Nash, 2000"] and the `multiplicative first_order` case is commonly known as beta correction [[40] "Haftka, 1991"]. For the `combined` correction, the use of both additive and multiplicative corrections allows the satisfaction of an additional matching condition, typically the truth function values at the previous correction point (e.g., the center of the previous trust region). The `combined` correction is then a multipoint

correction, as opposed to the local `additive` and `multiplicative` corrections. Each of these correction capabilities is described in detail in [[23] "Eldred et al., 2004"].

The correction factors force the surrogate models to match the true function values and possibly true function derivatives at the center point of each trust region. Currently, Dakota supports either zeroth-, first-, or second-order accurate correction methods, each of which can be applied using either an additive, multiplicative, or combined correction function. For each of these correction approaches, the correction is applied to the surrogate model and the corrected model is then interfaced with whatever algorithm is being employed. The default behavior is that no correction factor is applied.

The simplest correction approaches are those that enforce consistency in function values between the surrogate and original models at a single point in parameter space through use of a simple scalar offset or scaling applied to the surrogate model. First-order corrections such as the first-order multiplicative correction (also known as beta correction[14]) and the first-order additive correction [[58] "Lewis and Nash, 2000"] also enforce consistency in the gradients and provide a much more substantial correction capability that is sufficient for ensuring provable convergence in SBO algorithms. SBO convergence rates can be further accelerated through the use of second-order corrections which also enforce consistency in the Hessians[23], where the second-order information may involve analytic, finite-difference, or quasi-Newton Hessians.

Correcting surrogate models with additive corrections involves

$$\hat{f_{hi_\alpha}}(\mathbf{x}) = f_{lo}(\mathbf{x}) + \alpha(\mathbf{x}) \tag{5.1}$$

where multifidelity notation has been adopted for clarity. For multiplicative approaches, corrections take the form

$$\hat{f_{hi_\beta}}(\mathbf{x}) = f_{lo}(\mathbf{x})\beta(\mathbf{x}) \tag{5.2}$$

where, for local corrections, $\alpha(\mathbf{x})$ and $\beta(\mathbf{x})$ are first or second-order Taylor series approximations to the exact correction functions:

$$\alpha(\mathbf{x}) = A(\mathbf{x_c}) + \nabla A(\mathbf{x_c})^T(\mathbf{x}-\mathbf{x_c}) + \frac{1}{2}(\mathbf{x}-\mathbf{x_c})^T\nabla^2 A(\mathbf{x_c})(\mathbf{x}-\mathbf{x_c}) \tag{5.3}$$

$$\beta(\mathbf{x}) = B(\mathbf{x_c}) + \nabla B(\mathbf{x_c})^T(\mathbf{x}-\mathbf{x_c}) + \frac{1}{2}(\mathbf{x}-\mathbf{x_c})^T\nabla^2 B(\mathbf{x_c})(\mathbf{x}-\mathbf{x_c}) \tag{5.4}$$

where the exact correction functions are

$$A(\mathbf{x}) = f_{hi}(\mathbf{x}) - f_{lo}(\mathbf{x}) \tag{5.5}$$

$$B(\mathbf{x}) = \frac{f_{hi}(\mathbf{x})}{f_{lo}(\mathbf{x})} \tag{5.6}$$

Refer to[23] for additional details on the derivations.

A combination of additive and multiplicative corrections can provide for additional flexibility in minimizing the impact of the correction away from the trust region center. In other words, both additive and multiplicative corrections can satisfy local consistency, but through the combination, global accuracy can be addressed as well. This involves a convex combination of the additive and multiplicative corrections:

$$\hat{f_{hi_\gamma}}(\mathbf{x}) = \gamma\hat{f_{hi_\alpha}}(\mathbf{x}) + (1-\gamma)\hat{f_{hi_\beta}}(\mathbf{x})$$

where γ is calculated to satisfy an additional matching condition, such as matching values at the previous design iterate.

It should be noted that in both first order correction methods, the function $\hat{f}(x)$ matches the function value and gradients of $f_t(x)$ at $x = x_c$. This property is necessary in proving that the first order-corrected SBO algorithms are provably convergent to a local minimum of $f_t(x)$. However, the first order correction methods are significantly more expensive than the zeroth order correction methods, since the first order methods require computing both

$\nabla f_t(x_c)$ and $\nabla f_s(x_c)$. When the SBO strategy is used with either of the zeroth order correction methods, or with no correction method, convergence is not guaranteed to a local minimum of $f_t(x)$. That is, the SBO strategy becomes a heuristic optimization algorithm. From a mathematical point of view this is undesirable, but as a practical matter, the heuristic variants of SBO are often effective in finding local minima.

Usage guidelines

- Both the `additive` zeroth_order and `multiplicative` zeroth_order correction methods are "free" since they use values of $f_t(x_c)$ that are normally computed by the SBO strategy.
- The use of either the `additive` first_order method or the `multiplicative` first_order method does not necessarily improve the rate of convergence of the SBO algorithm.
- When using the first order correction methods, the `TRUE_FCN_GRAD` response keywords must be modified (see bottom of Figure to allow either analytic or numerical gradients to be computed. This provides the gradient data needed to compute the correction function.
- For many computationally expensive engineering optimization problems, gradients often are too expensive to obtain or are discontinuous (or may not exist at all). In such cases the heuristic SBO algorithm has been an effective approach at identifying optimal designs[35].

—>

zeroth_order

- Keywords Area
- model
- surrogate
- global
- correction
- zeroth_order

Specify that truth values must be matched.

Specification

Alias: none

Argument(s): none

Description

The correction specification specifies that the approximation will be corrected to match truth data. The keyword zeroth_order matching ensures that truth values are matched.

first_order

- Keywords Area
- model
- surrogate
- global
- correction
- first_order

Specify that truth values and gradients must be matched.

Specification

Alias: none
Argument(s): none

Description

The correction specification specifies that the approximation will be corrected to match truth data. The keyword zeroth_order matching ensures that truth values and gradients are matched.

second_order

- Keywords Area
- model
- surrogate
- global
- correction
- second_order

Specify that truth values, gradients and Hessians must be matched.

Specification

Alias: none
Argument(s): none

Description

The correction specification specifies that the approximation will be corrected to match truth data. The keyword zeroth_order matching ensures that truth values, gradients and Hessians are matched.

additive

- Keywords Area
- model
- surrogate
- global
- correction
- additive

Use an additive correction factor to improve the local accuracy of a surrogate.

Specification

Alias: none
Argument(s): none

Description

Use an additive correction factor to improve the local accuracy of a surrogate.

multiplicative

- Keywords Area
- model
- surrogate
- global
- correction
- multiplicative

Use a multiplicative correction factor to improve the local accuracy of a surrogatel.

Specification

Alias: none
Argument(s): none

Description

Use a multiplicative correction factor to improve the local accuracy of a surrogatel.

combined

- Keywords Area
- model
- surrogate
- global
- correction
- combined

A multipoint correction for a hierarchical surrogate.

Specification

Alias: none
Argument(s): none

Description

For the combined correction, the use of both additive and multiplicative corrections allows the satisfaction of an additional matching condition, typically the truth function values at the previous correction point (e.g., the center of the previous trust region). The combined correction is then a multipoint correction, as opposed to the local additive and multiplicative corrections

metrics

- Keywords Area
- model
- surrogate
- global
- metrics

Compute metrics of surrogate quality

Topics

This keyword is related to the topics:

- problem

Specification

Alias: diagnostics
Argument(s): STRINGLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		cross_validate	k-fold cross validation
	Optional		press	Leave-one-out cross-validation

Description

A variety of diagnostic metrics are available to assess the goodness of fit of a global surrogate to its training data.

The default diagnostics are:

- `sum_squared`
- `mean_squared`
- `root_mean_squared`
- `sum_abs`
- `mean_abs`
- `max_abs`
- `rsquared`

Additional metrics can be activated using the keywords `press` and `cross_validate`.

Theory

Most of these diagnostics refer to some operation on the residuals (the difference between the surrogate model and the truth model at the data points upon which the surrogate is built).

For example, `sum_squared` refers to the sum of the squared residuals, and `mean_abs` refers to the mean of the absolute value of the residuals. `rsquared` refers to the R-squared value typically used in regression analysis (the proportion of the variability in the response that can be accounted for by the surrogate model). Care should be taken when interpreting metrics, for example, errors may be near zero for interpolatory models or rsquared may not be applicable for non-polynomial models.

FAQ

Are the default diagnostics correct? Actually, this behavior is not in the ref man at all, so I just made this page up.

cross_validate

- Keywords Area
- model
- surrogate
- global
- metrics
- cross_validate

k-fold cross validation

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional **Optional***(Choose One)*	Description of Group **Group 1**	Dakota Keyword folds	Dakota Keyword Description Cross validation option
			percent	Cross validation option

Description

General k-fold cross validation may be performed by specifying `cross_validation`.
The cross-validation statistics will be calculated for all metrics.
Cross validation may include either

- `folds`, the number of folds into which to divide the build data (between 2 and number of data points) or
- `percent`, the fraction of data (between 0 and 0.5) to use in each fold.

These will be adjusted as needed based on the number of available training points.

FAQ

What happens if neither is chosen??

folds

- Keywords Area
- model
- surrogate
- global
- metrics
- cross_validate
- folds

Cross validation option

Specification

Alias: none
Argument(s): INTEGER

Description

See parent page

percent

- Keywords Area
- model
- surrogate
- global
- metrics
- cross_validate
- percent

Cross validation option

Specification

Alias: none
Argument(s): REAL

Description

See parent page

press

- Keywords Area
- model
- surrogate
- global
- metrics
- press

Leave-one-out cross-validation

Specification

Alias: none
Argument(s): none

Description

Leave-one-out (PRESS) cross validation may be performed by specifying `press`.
The cross-validation statistics will be calculated for all metrics.

challenge_points_file

- Keywords Area
- model
- surrogate
- global
- challenge_points_file

Specify a datafile of points to compare surrogate responses against

Specification

Alias: none

Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional(*Choose One*)	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

Description

If a file is specified, it must contain variables and response data.

The surrogate is evaluated at the points in the file, and the surrogate responses are compared against the results in the file.

The requested metrics will be reported for the surrogate evaluated at the specified variables versus the function values from the challenge data file.

The imported points file may be in annotated for free-form tabular format.

FAQ

What metrics?

annotated

- Keywords Area
- model
- surrogate
- global
- challenge_points_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- model
- surrogate
- global
- challenge_points_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

multipoint

- Keywords Area
- model
- surrogate
- multipoint

Construct a surrogate from multiple existing training points

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		tana	Local multi-point model via two-point nonlinear approximation
	Required		actual_model_-pointer	Pointer to specify a "truth" model, from which to construct a surrogate

Description

Multipoint approximations use data from previous design points to improve the accuracy of local approximations. The data often comes from the current and previous iterates of a minimization algorithm.

Currently, only the Two-point Adaptive Nonlinearity Approximation (TANA-3) method of [[92] "Xu and Grandhi, 1998"] is supported with the `tana` keyword.

The truth model to be used to generate the value/gradient data used in the approximation is identified through the required `actual_model_pointer` specification.

See Also

These keywords may also be of interest:

- local
- global
- hierarchical

tana

- Keywords Area
- model
- surrogate
- multipoint
- tana

Local multi-point model via two-point nonlinear approximation

Specification

Alias: none

Argument(s): none

Description

TANA stands for Two Point Adaptive Nonlinearity Approximation.

The TANA-3 method [[92] "Xu and Grandhi, 1998"] is a multipoint approximation method based on the two point exponential approximation[24]. This approach involves a Taylor series approximation in intermediate variables where the powers used for the intermediate variables are selected to match information at the current and previous expansion points.

Theory

The form of the TANA model is:

$$\hat{f}(\mathbf{x}) \approx f(\mathbf{x}_2) + \sum_{i=1}^{n} \frac{\partial f}{\partial x_i}(\mathbf{x}_2) \frac{x_{i,2}^{1-p_i}}{p_i}(x_i^{p_i} - x_{i,2}^{p_i}) + \frac{1}{2}\epsilon(\mathbf{x}) \sum_{i=1}^{n} (x_i^{p_i} - x_{i,2}^{p_i})^2$$

where n is the number of variables and:

$$p_i = 1 + \ln\left[\frac{\frac{\partial f}{\partial x_i}(\mathbf{x}_1)}{\frac{\partial f}{\partial x_i}(\mathbf{x}_2)}\right] \Big/ \ln\left[\frac{x_{i,1}}{x_{i,2}}\right] \quad \epsilon(\mathbf{x}) = \frac{H}{\sum_{i=1}^{n}(x_i^{p_i} - x_{i,1}^{p_i})^2 + \sum_{i=1}^{n}(x_i^{p_i} - x_{i,2}^{p_i})^2} \quad H = 2\left[f(\mathbf{x}_1) - f(\mathbf{x}_2) - \sum_{i=1}^{n} \frac{\partial f}{\partial x_i}(\mathbf{x}_2) \frac{x_{i,2}^{1-p}}{p_i}\right.$$

and $\mathbf{x}_2$ and $\mathbf{x}_1$ are the current and previous expansion points. Prior to the availability of two expansion points, a first-order Taylor series is used.

actual_model_pointer

- Keywords Area
- model
- surrogate
- multipoint
- actual_model_pointer

Pointer to specify a "truth" model, from which to construct a surrogate

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

This must point to a model block, identified by id_model. That model will be run to generate training data, from which a surrogate model will be constructed.

See block_pointer for details about pointers.

local

- Keywords Area
- model
- surrogate
- local

Build a locally accurate surrogate from data at a single point

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		taylor_series	Construct a Taylor Series expansion around a point
	Required		actual_model_-pointer	Pointer to specify a "truth" model, from which to construct a surrogate

Description

Local approximations use value, gradient, and possibly Hessian data from a single point to form a series expansion for approximating data in the vicinity of this point.

The currently available local approximation is the `taylor_series` selection.

The truth model to be used to generate the value/gradient/Hessian data used in the series expansion is identified through the required `actual_model_pointer` specification. The use of a model pointer (as opposed to an interface pointer) allows additional flexibility in defining the approximation. In particular, the derivative specification for the truth model may differ from the derivative specification for the approximation, and the truth model results being approximated may involve a model recursion (e.g., the values/gradients from a nested model).

See Also

These keywords may also be of interest:

- global
- hierarchical
- multipoint

taylor_series

- Keywords Area
- model
- surrogate
- local
- taylor_series

Construct a Taylor Series expansion around a point

Specification

Alias: none

Argument(s): none

Description

The Taylor series model is purely a local approximation method. That is, it provides local trends in the vicinity of a single point in parameter space.

The order of the Taylor series may be either first-order or second-order, which is automatically determined from the gradient and Hessian specifications in the responses specification (see responses for info on how to specify gradients and Hessians) for the truth model.

Theory

The first-order Taylor series expansion is:

$$\hat{f}(\mathbf{x}) \approx f(\mathbf{x}_0) + \nabla_{\mathbf{x}} f(\mathbf{x}_0)^T (\mathbf{x} - \mathbf{x}_0) \tag{5.7}$$

and the second-order expansion is:

$$\hat{f}(\mathbf{x}) \approx f(\mathbf{x}_0) + \nabla_{\mathbf{x}} f(\mathbf{x}_0)^T (\mathbf{x} - \mathbf{x}_0) + \frac{1}{2}(\mathbf{x} - \mathbf{x}_0)^T \nabla^2_{\mathbf{x}} f(\mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0) \tag{5.8}$$

where $\mathbf{x}_0$ is the expansion point in n-dimensional parameter space and $f(\mathbf{x}_0)$, $\nabla_{\mathbf{x}} f(\mathbf{x}_0)$, and $\nabla^2_{\mathbf{x}} f(\mathbf{x}_0)$ are the computed response value, gradient, and Hessian at the expansion point, respectively.

As dictated by the responses specification used in building the local surrogate, the gradient may be analytic or numerical and the Hessian may be analytic, numerical, or based on quasi-Newton secant updates.

In general, the Taylor series model is accurate only in the region of parameter space that is close to $\mathbf{x}_0$. While the accuracy is limited, the first-order Taylor series model reproduces the correct value and gradient at the point $\mathbf{x}_0$, and the second-order Taylor series model reproduces the correct value, gradient, and Hessian. This consistency is useful in provably-convergent surrogate-based optimization. The other surface fitting methods do not use gradient information directly in their models, and these methods rely on an external correction procedure in order to satisfy the consistency requirements of provably-convergent SBO.

actual_model_pointer

- Keywords Area
- model
- surrogate
- local
- actual_model_pointer

Pointer to specify a ”truth” model, from which to construct a surrogate

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none

Argument(s): STRING

Description

This must point to a model block, identified by id_model. That model will be run to generate training data, from which a surrogate model will be constructed.

See block_pointer for details about pointers.

hierarchical

- Keywords Area
- model
- surrogate
- hierarchical

Hierarchical approximations use corrected results from a low fidelity model as an approximation to the results of a high fidelity "truth" model.

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		low_fidelity_-model_pointer	
	Required		high_fidelity_-model_pointer	
	Required		correction	Correction approaches for surrogate models

Description

Hierarchical approximations use corrected results from a low fidelity model as an approximation to the results of a high fidelity "truth" model. These approximations are also known as model hierarchy, multifidelity, variable fidelity, and variable complexity approximations. The required `low_fidelity_model_pointer` specification points to the low fidelity model specification. This model is used to generate low fidelity responses which are then corrected and returned to an iterator. The required `high_fidelity_model_pointer` specification points to the specification for the high fidelity truth model. This model is used only for verifying low fidelity results and updating low fidelity corrections. The `correction` specification specifies which correction technique will be applied to the low fidelity results in order to match the high fidelity results at one or more points. In the hierarchical case (as compared to the global case), the `correction` specification is required, since the omission of a correction technique would effectively eliminate the purpose of the high fidelity model. If it is desired to use a low fidelity model without corrections, then a hierarchical approximation is not needed and a `single` model should be used. Refer to global for additional information on available correction approaches.

Theory

Multifidelity Surrogates : Multifidelity modeling involves the use of a low-fidelity physics-based model as a surrogate for the original high-fidelity model. The low-fidelity model typically involves a coarser mesh, looser convergence tolerances, reduced element order, or omitted physics. It is a separate model in its own right and does not require data from the high-fidelity model for construction. Rather, the primary need for high-fidelity evaluations is for defining correction functions that are applied to the low-fidelity results.

Multifidelity Surrogate Models

A second type of surrogate is the {*model* hierarchy} type (also called multifidelity, variable fidelity, variable complexity, etc.). In this case, a model that is still physics-based but is of lower fidelity (e.g., coarser discretization, reduced element order, looser convergence tolerances, omitted physics) is used as the surrogate in place of the high-fidelity model. For example, an inviscid, incompressible Euler CFD model on a coarse discretization could be used as a low-fidelity surrogate for a high-fidelity Navier-Stokes model on a fine discretization.

See Also

These keywords may also be of interest:

- global
- local
- multipoint

low_fidelity_model_pointer

- Keywords Area
- model
- surrogate
- hierarchical
- low_fidelity_model_pointer

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none

Argument(s): STRING

high_fidelity_model_pointer

- Keywords Area
- model
- surrogate
- hierarchical
- high_fidelity_model_pointer

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): STRING

correction

- Keywords Area
- model
- surrogate
- hierarchical
- correction

Correction approaches for surrogate models

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	zeroth_order	Specify that truth values must be matched.

			first_order	Specify that truth values and gradients must be matched.
			second_order	Specify that truth values, gradients and Hessians must be matched.
	Required*(Choose One)*	**Group 2**	additive	Use an additive correction factor to improve the local accuracy of a surrogate.
			multiplicative	Use a multiplicative correction factor to improve the local accuracy of a surrogatel.
			combined	A multipoint correction for a hierarchical surrogate.

Description

Some of the surrogate model types support the use of correction factors that improve the local accuracy of the surrogate models.

The `correction` specification specifies that the approximation will be corrected to match truth data, either matching truth values in the case of `zeroth_order` matching, matching truth values and gradients in the case of `first_order` matching, or matching truth values, gradients, and Hessians in the case of `second_order` matching. For `additive` and `multiplicative` corrections, the correction is local in that the truth data is matched at a single point, typically the center of the approximation region. The `additive` correction adds a scalar offset (`zeroth_order`), a linear function (`first_order`), or a quadratic function (`second_order`) to the approximation to match the truth data at the point, and the `multiplicative` correction multiplies the approximation by a scalar (`zeroth_order`), a linear function (`first_order`), or a quadratic function (`second_order`) to match the truth data at the point. The `additive first_order` case is due to [[58] "-Lewis and Nash, 2000"] and the `multiplicative first_order` case is commonly known as beta correction [[40] "Haftka, 1991"]. For the `combined` correction, the use of both additive and multiplicative corrections allows the satisfaction of an additional matching condition, typically the truth function values at the previous correction point (e.g., the center of the previous trust region). The `combined` correction is then a multipoint correction, as opposed to the local `additive` and `multiplicative` corrections. Each of these correction capabilities is described in detail in [[23] "Eldred et al., 2004"].

The correction factors force the surrogate models to match the true function values and possibly true function derivatives at the center point of each trust region. Currently, Dakota supports either zeroth-, first-, or second-order accurate correction methods, each of which can be applied using either an additive, multiplicative, or combined correction function. For each of these correction approaches, the correction is applied to the surrogate model and the corrected model is then interfaced with whatever algorithm is being employed. The default behavior is that no correction factor is applied.

The simplest correction approaches are those that enforce consistency in function values between the surrogate

and original models at a single point in parameter space through use of a simple scalar offset or scaling applied to the surrogate model. First-order corrections such as the first-order multiplicative correction (also known as beta correction[14]) and the first-order additive correction [[58] "Lewis and Nash, 2000"] also enforce consistency in the gradients and provide a much more substantial correction capability that is sufficient for ensuring provable convergence in SBO algorithms. SBO convergence rates can be further accelerated through the use of second-order corrections which also enforce consistency in the Hessians[23], where the second-order information may involve analytic, finite-difference, or quasi-Newton Hessians.

Correcting surrogate models with additive corrections involves

$$\hat{f_{hi_\alpha}}(\mathbf{x}) = f_{lo}(\mathbf{x}) + \alpha(\mathbf{x}) \tag{5.9}$$

where multifidelity notation has been adopted for clarity. For multiplicative approaches, corrections take the form

$$\hat{f_{hi_\beta}}(\mathbf{x}) = f_{lo}(\mathbf{x})\beta(\mathbf{x}) \tag{5.10}$$

where, for local corrections, $\alpha(\mathbf{x})$ and $\beta(\mathbf{x})$ are first or second-order Taylor series approximations to the exact correction functions:

$$\alpha(\mathbf{x}) = A(\mathbf{x_c}) + \nabla A(\mathbf{x_c})^T(\mathbf{x}-\mathbf{x_c}) + \frac{1}{2}(\mathbf{x}-\mathbf{x_c})^T\nabla^2 A(\mathbf{x_c})(\mathbf{x}-\mathbf{x_c}) \tag{5.11}$$

$$\beta(\mathbf{x}) = B(\mathbf{x_c}) + \nabla B(\mathbf{x_c})^T(\mathbf{x}-\mathbf{x_c}) + \frac{1}{2}(\mathbf{x}-\mathbf{x_c})^T\nabla^2 B(\mathbf{x_c})(\mathbf{x}-\mathbf{x_c}) \tag{5.12}$$

where the exact correction functions are

$$A(\mathbf{x}) = f_{hi}(\mathbf{x}) - f_{lo}(\mathbf{x}) \tag{5.13}$$

$$B(\mathbf{x}) = \frac{f_{hi}(\mathbf{x})}{f_{lo}(\mathbf{x})} \tag{5.14}$$

Refer to[23] for additional details on the derivations.

A combination of additive and multiplicative corrections can provide for additional flexibility in minimizing the impact of the correction away from the trust region center. In other words, both additive and multiplicative corrections can satisfy local consistency, but through the combination, global accuracy can be addressed as well. This involves a convex combination of the additive and multiplicative corrections:

$$\hat{f_{hi_\gamma}}(\mathbf{x}) = \gamma\hat{f_{hi_\alpha}}(\mathbf{x}) + (1-\gamma)\hat{f_{hi_\beta}}(\mathbf{x})$$

where γ is calculated to satisfy an additional matching condition, such as matching values at the previous design iterate.

It should be noted that in both first order correction methods, the function $\hat{f}(x)$ matches the function value and gradients of $f_t(x)$ at $x = x_c$. This property is necessary in proving that the first order-corrected SBO algorithms are provably convergent to a local minimum of $f_t(x)$. However, the first order correction methods are significantly more expensive than the zeroth order correction methods, since the first order methods require computing both $\nabla f_t(x_c)$ and $\nabla f_s(x_c)$. When the SBO strategy is used with either of the zeroth order correction methods, or with no correction method, convergence is not guaranteed to a local minimum of $f_t(x)$. That is, the SBO strategy becomes a heuristic optimization algorithm. From a mathematical point of view this is undesirable, but as a practical matter, the heuristic variants of SBO are often effective in finding local minima.

Usage guidelines

- Both the `additive` zeroth_order and `multiplicative` zeroth_order correction methods are "free" since they use values of $f_t(x_c)$ that are normally computed by the SBO strategy.
- The use of either the `additive` first_order method or the `multiplicative` first_order method does not necessarily improve the rate of convergence of the SBO algorithm.

- When using the first order correction methods, the `TRUE_FCN_GRAD` response keywords must be modified (see bottom of Figure to allow either analytic or numerical gradients to be computed. This provides the gradient data needed to compute the correction function.
- For many computationally expensive engineering optimization problems, gradients often are too expensive to obtain or are discontinuous (or may not exist at all). In such cases the heuristic SBO algorithm has been an effective approach at identifying optimal designs[35].

—>

zeroth_order

- Keywords Area
- model
- surrogate
- hierarchical
- correction
- zeroth_order

Specify that truth values must be matched.

Specification

Alias: none
Argument(s): none

Description

The correction specification specifies that the approximation will be corrected to match truth data. The keyword zeroth_order matching ensures that truth values are matched.

first_order

- Keywords Area
- model
- surrogate
- hierarchical
- correction
- first_order

Specify that truth values and gradients must be matched.

Specification

Alias: none
Argument(s): none

Description

The correction specification specifies that the approximation will be corrected to match truth data. The keyword zeroth_order matching ensures that truth values and gradients are matched.

second_order

- Keywords Area
- model
- surrogate
- hierarchical
- correction
- second_order

Specify that truth values, gradients and Hessians must be matched.

Specification

Alias: none
Argument(s): none

Description

The correction specification specifies that the approximation will be corrected to match truth data. The keyword zeroth_order matching ensures that truth values, gradients and Hessians are matched.

additive

- Keywords Area
- model
- surrogate
- hierarchical
- correction
- additive

Use an additive correction factor to improve the local accuracy of a surrogate.

Specification

Alias: none
Argument(s): none

Description

Use an additive correction factor to improve the local accuracy of a surrogate.

multiplicative

- Keywords Area
- model
- surrogate
- hierarchical
- correction
- multiplicative

Use a multiplicative correction factor to improve the local accuracy of a surrogatel.

Specification

Alias: none
Argument(s): none

Description

Use a multiplicative correction factor to improve the local accuracy of a surrogatel.

combined

- Keywords Area
- model
- surrogate
- hierarchical
- correction
- combined

A multipoint correction for a hierarchical surrogate.

Specification

Alias: none
Argument(s): none

Description

For the combined correction, the use of both additive and multiplicative corrections allows the satisfaction of an additional matching condition, typically the truth function values at the previous correction point (e.g., the center of the previous trust region). The combined correction is then a multipoint correction, as opposed to the local additive and multiplicative corrections

5.3.7 nested

- Keywords Area
- model
- nested

A model whose responses are computed through the use of a sub-iterator

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		optional_interface-_pointer sub_method_-pointer	The `sub_method_-pointer` specifies the method block for the sub-iterator
	Required			

Description

Instead of appealing directly to a primary interface, a nested model maps variables to responses by executing a secondary iterator, or a "sub-iterator". In other words, a function evaluation of the primary study consists of a solution of an entire secondary study - potentially many secondary function evaluations.

The sub-iterator in turn operates on a sub-model. The sub-iterator responses may be combined with non-nested contributions from an optional interface specification.

A **`sub_method_pointer`** must be provided in order to specify the method block describing the sub-iterator. The remainder of the model is specified under that keyword.

A **`optional_interface_pointer` points** to the interface specification and `optional_interface-_responses_pointer` points to a responses specification describing the data to be returned by this interface). This interface is used to provide non-nested data, which is then combined with data from the nested iterator using the `primary_response_mapping` and `secondary_response_mapping` inputs (see mapping discussion below).

Examples

An example of variable and response mappings is provided below:

```
primary_variable_mapping  = '' '' 'X'   'Y'
secondary_variable_mapping = '' '' 'mean' 'mean'
primary_response_mapping  = 1. 0. 0. 0. 0. 0. 0. 0. 0.
secondary_response_mapping = 0. 0. 0. 1. 3. 0. 0. 0. 0.
               0. 0. 0. 0. 0. 0. 1. 3. 0.
```

The variable mappings correspond to 4 top-level variables, the first two of which employ the default mappings from active top-level variables to sub-model variables of the same type (option 3 above) and the latter two of which are inserted into the mean distribution parameters of sub-model variables `'X'` and `'Y'` (option 1 above). The response mappings define a 3 by 9 matrix corresponding to 9 inner loop response attributes and 3 outer loop response functions (one primary response function and 2 secondary functions, such as one objective and two constraints). Each row of the response mapping is a vector which is multiplied (i.e, with a dot-product) against the 9 sub-iterator values to determine the outer loop function. Consider a UQ example with 3 response functions, each providing a mean, a standard deviation, and one level mapping (if no level mappings are specified, the responses would only have a mean and standard deviation). The primary response mapping can be seen to extract the first value from the inner loop, which would correspond to the mean of the first response function. This mapped sub-iterator response becomes a single objective function, least squares term, or generic response function at the outer level, as dictated by the top-level response specification. The secondary response mapping maps the fourth sub-iterator response function plus 3 times the fifth sub-iterator response function (mean plus 3 standard deviations) into one top-level nonlinear constraint and the seventh sub-iterator response function plus 3 times the eighth sub-iterator response function (mean plus 3 standard deviations) into another top-level nonlinear constraint, where these top-level nonlinear constraints may be inequality or equality, as dictated by the top-level response specification. Note that a common case is for each sub-iterator response to be mapped to a unique outer loop response (for example, in the nested UQ case where one wants to determine an interval on each inner loop statistic). In these simple cases, the response mapping would define an identity matrix.

Theory

problem

See Also

These keywords may also be of interest:

- single
- surrogate

optional_interface_pointer

- Keywords Area
- model
- nested
- optional_interface_pointer

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- block_pointer
- empty

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		optional_interface-_responses_pointer	

optional_interface_responses_pointer

- Keywords Area
- model
- nested
- optional_interface_pointer
- optional_interface_responses_pointer

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- block_pointer
- empty

Specification

Alias: none
Argument(s): STRING

sub_method_pointer

- Keywords Area
- model
- nested
- sub_method_pointer

The `sub_method_pointer` specifies the method block for the sub-iterator

Topics

This keyword is related to the topics:

- block_pointer
- not_yet_reviewed

Specification

Alias: none

Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		iterator_servers	Specify the number of iterator servers when Dakota is run in parallel
	Optional		iterator_scheduling	Specify the scheduling of concurrent iterators when Dakota is run in parallel
	Optional		processors_per_-iterator	Specify the number of processors per iterator server when Dakota is run in parallel
	Optional		primary_variable_-mapping	
	Optional		secondary_-variable_mapping	
	Optional		primary_response_-mapping	
	Optional		secondary_-response_mapping	

Description

The `sub_method_pointer` specifies the method block for the sub-iterator.

See block_pointer for details about pointers.

Nested models may employ mappings for both the variable inputs to the sub-model and the response outputs from the sub-model. In the former case, the `primary_variable_mapping` and `secondary_variable_mapping` specifications are used to map from the active top-level variables into the sub-model variables, and in the latter case, the `primary_response_mapping` and `secondary_response_mapping` specifications are used to compute the sub-model response contributions to the top-level responses.

For the variable mappings, the primary and secondary specifications provide lists of strings which are used to target specific sub-model variables and their sub-parameters, respectively. The primary strings are matched to continuous or discrete variable labels such as `'cdv_1'` (either user-supplied or default labels), and the secondary

strings are matched to either real or integer random variable distribution parameters such as `'mean'` or `'num_trials'` (the form of the uncertain distribution parameter keyword that is appropriate for a single variable instance) or continuous or discrete design/state variable sub-parameters such as `'lower_bound'` or `'upper_bound'` (again, keyword form appropriate for a single variable instance). No coercion of types is supported, so real-valued top-level variables should map to either real-valued sub-model variables or real-valued sub-parameters and integer-valued top-level variables should map to either integer-valued sub-model variables or integer-valued sub-parameters. As long as these real versus integer constraints are satisfied, mappings are free to cross variable types (design, aleatory uncertain, epistemic uncertain, state) and domain types (continuous, discrete). Both `primary_variable_mapping` and `secondary_variable_mapping` specifications are optional, which is designed to support the following three possibilities:

1. If both primary and secondary variable mappings are specified, then an active top-level variable value will be inserted into the identified sub-parameter (the secondary mapping) for the identified sub-model variable (the primary mapping).

2. If a primary mapping is specified but a secondary mapping is not, then an active top-level variable value will be inserted into the identified sub-model variable value (the primary mapping).

3. If a primary mapping is not specified (corresponding secondary mappings, if specified, are ignored), then an active top-level variable value will be inserted into a corresponding sub-model variable, based on matching of variable types (e.g., top-level and sub-model variable specifications both allocate a set of `'continuous_design'` variables which are active at the top level). Multiple sub-model variable types may be updated in this manner, provided that they are all active in the top-level variables. Since there is a direct variable correspondence for these default insertions, sub-model bounds and labels are also updated from the top-level bounds and labels in order to eliminate the need for redundant input file specifications. Thus, it is typical for the sub-model variables specification to only contain the minimal required information, such as the number of variables, for these insertion targets. The sub-model must allocate enough space for each of the types that will accept default insertions, and the leading set of matching sub-model variables are updated (i.e., the sub-model may allocate more than needed and the trailing set will be unmodified).

These different variable mapping possibilities may be used in any combination by employing empty strings (`''`) for particular omitted mappings (the number of strings in user-supplied primary and secondary variable mapping specifications must equal the total number of active top-level variables, including both continuous and discrete types). The ordering of the active variables is the same as shown in dakota.input.summary on Input Spec Summary and as presented in variables.

If inactive variables are present at the outer level, then the default type 3 mapping is used for these variables; that is, outer loop inactive variables are inserted into inner loop variables (active or inactive) based on matching of variable types, top-level bounds and labels are also propagated, the inner loop must allocate sufficient space to receive the outer loop values, and the leading subset within this inner loop allocation is updated. This capability is important for allowing nesting beyond two levels, since an active variable at the outer-most loop may become inactive at the next lower level, but still needs to be further propagated down to lower levels in the recursion.

For the response mappings, the primary and secondary specifications provide real-valued multipliers to be applied to sub-iterator response results so that the responses from the inner loop can be mapped into a new set of responses at the outer loop. For example, if the nested model is being employed within a mixed aleatory-epistemic uncertainty quantification, then aleatory statistics from the inner loop (such as moments of the response) are mapped to the outer level, where minima and maxima of these aleatory statistics are computed as functions of the epistemic parameters. The response mapping defines a matrix which scales the values from the inner loop and determines their position in the outer loop response vector. Each row of the mapping corresponds to one outer loop response, where each column of the mapping corresponds to a value from the inner loop. Depending on the number of responses and the particular attributes calculated on the inner loop, there will be a vector of inner loop response values that need to be accounted for in the mapping. This vector of inner loop response results is defined as follows for different sub-iterator types:

- optimization: the final objective function(s) and nonlinear constraints
- nonlinear least squares: the final least squares terms and nonlinear constraints
- aleatory uncertainty quantification (UQ): for each response function, a mean statistic, a standard deviation statistic, and all probability/reliability/generalized reliability/response level results for any user-specified `response_levels`, `probability_levels`, `reliability_levels`, and/or `gen_reliability_levels`, in that order.
- epistemic and mixed aleatory/epistemic UQ using interval estimation methods: lower and upper interval bounds for each response function.
- epistemic and mixed aleatory/epistemic UQ using evidence methods: for each response function, lower and upper interval bounds (belief and plausibility) for all probability/reliability/generalized reliability/response level results computed from any user-specified `response_levels`, `probability_levels`, `reliability_levels`, and/or `gen_reliability_levels`, in that order.
- parameter studies and design of experiments: for optimization and least squares response data sets, the best solution found (lowest constraint violation if infeasible, lowest composite objective function if feasible). For generic response data sets, a best solution metric is not defined, so the sub-iterator response vector is empty in this case.

The primary values map sub-iterator response results into top-level objective functions, least squares terms, or generic response functions, depending on the declared top-level response set. The secondary values map sub-iterator response results into top-level nonlinear inequality and equality constraints.

Nested models utilize a sub-iterator and a sub-model to perform a complete iterative study as part of every evaluation of the model. This sub-iteration accepts variables from the outer level, performs the sub-level analysis, and computes a set of sub-level responses which are passed back up to the outer level. Mappings are employed for both the variable inputs to the sub-model and the response outputs from the sub-model.

In the variable mapping case, primary and secondary variable mapping specifications are used to map from the top-level variables into the sub-model variables. These mappings support three possibilities in any combination: (1) insertion of an active top-level variable value into an identified sub-model distribution parameter for an identified active sub-model variable, (2) insertion of an active top-level variable value into an identified active sub-model variable value, and (3) addition of an active top-level variable value as an inactive sub-model variable, augmenting the active sub-model variables.

In the response mapping case, primary and secondary response mapping specifications are used to map from the sub-model responses back to the top-level responses. These specifications provide real-valued multipliers that are applied to the sub-iterator response results to define the outer level response set. These nested data results may be combined with non-nested data through use of the "optional interface" component within nested models.

The nested model constructs admit a wide variety of multi-iterator, multi-model solution approaches. For example, optimization within optimization (for hierarchical multidisciplinary optimization), uncertainty quantification within uncertainty quantification (for second-order probability), uncertainty quantification within optimization (for optimization under uncertainty), and optimization within uncertainty quantification (for uncertainty of optima) are all supported, with and without surrogate model indirection. Several examples of nested model usage are provided in the Users Manual, most notably mixed epistemic-aleatory UQ, optimization under uncertainty (OUU), and surrogate-based UQ.

iterator_servers

- Keywords Area
- model

- nested
- sub_method_pointer
- iterator_servers

Specify the number of iterator servers when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `iterator_servers` specification supports user override of the automatic parallel configuration for the number of iterator servers. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired number of partitions at the iterator parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

iterator_scheduling

- Keywords Area
- model
- nested
- sub_method_pointer
- iterator_scheduling

Specify the scheduling of concurrent iterators when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required(*Choose One*)	**Group 1**	master	Specify a dedicated master partition for parallel iterator scheduling
			peer	Specify a peer partition for parallel iterator scheduling

Description

When Dakota is run in parallel, the partition type of the iterator servers is determined automatically. If this setting is undesirable, it may be overridden by the user using the `iterator_scheduling` keyword.

master

- Keywords Area
- model
- nested
- sub_method_pointer
- iterator_scheduling
- master

Specify a dedicated master partition for parallel iterator scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a dedicated master partition. In a dedicated master partition, one processor (the ”master”) dynamically schedules work on the iterator servers. This reduces the number of processors available to create servers by 1.

peer

- Keywords Area
- model
- nested
- sub_method_pointer
- iterator_scheduling
- peer

Specify a peer partition for parallel iterator scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a peer partition. In a peer partition, all processors are available to be assigned to iterator servers. Note that unlike the case of `evaluation_scheduling`, it is not possible to specify `static` or `dynamic`.

processors_per_iterator

- Keywords Area
- model
- nested
- sub_method_pointer
- processors_per_iterator

Specify the number of processors per iterator server when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `processors_per_iterator` specification supports user override of the automatic parallel configuration for the number of processors in each iterator server. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired server size at the iterator parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

primary_variable_mapping

- Keywords Area
- model
- nested
- sub_method_pointer
- primary_variable_mapping

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): STRINGLIST

secondary_variable_mapping

- Keywords Area
- model
- nested
- sub_method_pointer
- secondary_variable_mapping

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): STRINGLIST

primary_response_mapping

- Keywords Area
- model
- nested
- sub_method_pointer
- primary_response_mapping

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): REALLIST

secondary_response_mapping

- Keywords Area
- model
- nested
- sub_method_pointer
- secondary_response_mapping

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: none
Argument(s): REALLIST

5.4 variables

- Keywords Area
- variables

Specifies the parameter set to be iterated by a particular method.

Topics

This keyword is related to the topics:

- block

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		id_variables	Give the block a name, in case of multiple variables blocks
	Optional		active	Modify how the iterator treats different types of variables
	Optional(*Choose One*)	**Group 1**	mixed	Continuous and discrete variables are treated separately
			relaxed	Treat discrete variables as continuous
	Optional		continuous_design	Continuous design variables, each defined by a single interval
	Optional		discrete_design_-range	Discrete design variable - specified using a range of integers

	Optional	discrete_design_-set_integer	Discrete design variables - specified using lists of integers
	Optional	discrete_design_-set_real	Discrete design variables - specified using lists of reals
	Optional	normal_uncertain	Aleatory uncertain variable - normal (Gaussian)
	Optional	lognormal_-uncertain	Aleatory uncertain variable - lognormal
	Optional	uniform_uncertain	Aleatory uncertain variable - uniform
	Optional	loguniform_-uncertain	Aleatory uncertain variable - loguniform
	Optional	triangular_-uncertain	Aleatory uncertain variable - triangular
	Optional	exponential_-uncertain	Aleatory uncertain variable - exponential
	Optional	beta_uncertain	Aleatory uncertain variable - beta
	Optional	gamma_uncertain	Aleatory uncertain variable - gamma
	Optional	gumbel_uncertain	Aleatory uncertain variable - gumbel
	Optional	frechet_uncertain	Aleatory uncertain variable - Frechet
	Optional	weibull_uncertain	Aleatory uncertain variable - Weibull
	Optional	histogram_bin_-uncertain	Aleatory uncertain variable - an arbitrary continuous random variable !!!!!!!!!!!!!!clarify difference between ordinatea and count

	Optional	poisson_uncertain	Aleatory uncertain discrete variable - Poisson
	Optional	binomial_uncertain	Aleatory uncertain discrete variable - binomial
	Optional	negative_binomial-_uncertain	Aleatory uncertain discrete variable - negative binomial
	Optional	geometric_-uncertain	Aleatory uncertain discrete variable - geometric
	Optional	hypergeometric_-uncertain	Aleatory uncertain discrete variable - hypergeometric
	Optional	histogram_point_-uncertain	Aleatory uncertain discrete variable - arbitrary discrete random variable
	Optional	uncertain_-correlation_matrix	Specify correlations between aleatory uncertain variables
	Optional	continuous_-interval_uncertain	Epistemic uncertain variable - continuous values within one or more intervals
	Optional	discrete_interval_-uncertain	Epistemic uncertain variable - one or more intervals
	Optional	discrete_uncertain-_set_integer	Discrete, epistemic uncertain variable - integers within a set
	Optional	discrete_uncertain-_set_real	Discrete, epistemic uncertain variable - real numbers within a set
	Optional	continuous_state	Continuous state variables
	Optional	discrete_state_-range	Discrete state variable defined by bounds

	Optional	discrete_state_set_-integer	Discrete state variables, each defined by a set of permissible integers
	Optional	discrete_state_set_-real	Discrete state variables, each defined by a set of permissible real numbers

Description

The `variables` specification in a Dakota input file specifies the parameter set to be iterated by a particular method.

In the case of

- An optimization study:
  - These variables are adjusted in order to locate an optimal design.
- Parameter studies/sensitivity analysis/design of experiments:
  - These parameters are perturbed to explore the parameter space.
- Uncertainty analysis:
  - The variables are associated with distribution/interval characterizations which are used to compute corresponding distribution/interval characterizations for response functions.

To accommodate these different studies, Dakota supports different:

- Variable types
  - design
  - aleatory uncertain
  - epistemic uncertain
  - state
- Variable domains
  - continuous
  - discrete
    - discrete range
    - discrete integer set
    - discrete real set

Use the variables_and_responses page to browse the available variables by type and domain.

Variable Types

- Design Variables

  - Design variables are those variables which are modified for the purposes of computing an optimal design.
  - The most common type of design variables encountered in engineering applications are of the continuous type. These variables may assume any real value within their bounds.
  - All but a handful of the optimization algorithms in Dakota support continuous design variables exclusively.

- Aleatory Uncertain Variables
  - Aleatory uncertainty is also known as inherent variability, irreducible uncertainty, or randomness.
  - Aleatory uncertainty is predominantly characterized using probability theory. This is the only option implemented in Dakota.

- Epistemic Uncertain Variables
  - Epistemic uncertainty is uncertainty due to lack of knowledge.
  - In Dakota, epistemic uncertainty is characterized by interval analysis or the Dempster-Shafer theory of evidence.
  - Note that epistemic uncertainty can also be modeled with probability density functions (as done with aleatory uncertainty). Dakota does not support this capability.

- State Variables
  - State variables consist of "other" variables which are to be mapped through the simulation interface, in that they are not to be used for design and they are not modeled as being uncertain.
  - State variables provide a convenient mechanism for managing additional model parameterizations such as mesh density, simulation convergence tolerances, and time step controls.
  - Only parameter studies and design of experiments methods will iterate on state variables.
  - The `initial_value` is used as the only value for the state variable for all other methods, unless `active state` is invoked.
  - See more details on the state_variables page.

Variable Domains

Continuous variables are defined by bounds. Discrete variables can be defined in one of three ways, which are discussed on the page discrete_variables.

Ordering of Variables

The ordering of variables is important, and a consistent ordering is employed throughout the Dakota software. It is the same ordering as shown in dakota.input.summary That ordering can be summarized as

1. design
   (a) continuous
   (b) discrete integer
   (c) discrete real
2. aleatory uncertain
   (a) continuous
   (b) discrete integer

   (c) discrete real

3. epistemic uncertain
   (a) continuous
   (b) discrete integer
   (c) discrete real

4. state
   (a) continuous
   (b) discrete integer
   (c) discrete real

Ordering of variable types below this granularity (e.g., from normal to histogram bin within aleatory uncertain - continuous) is defined somewhat arbitrarily, but is enforced consistently throughout the code.

Active Variables

The reason variable types exist is that methods have the capability to treat variable types differently. All methods have default behavior that determines which variable types are "active" and will be assigned values by the method. For example, optimization methods will only vary the design variables - by default.

The default behavior should be described on each method page, or on topics pages that relate to classes of methods. In addition, the default behavior can be modified using the active keyword.

Examples

Several examples follow. In the first example, two continuous design variables are specified:

```
variables,
    continuous_design = 2
      initial_point    0.9    1.1
      upper_bounds     5.8    2.9
      lower_bounds     0.5   -2.9
      descriptors   'radius' 'location'
```

In the next example, defaults are employed. In this case, `initial_point` will default to a vector of `0.` values, `upper_bounds` will default to vector values of `DBL_MAX` (the maximum number representable in double precision for a particular platform, as defined in the platform's `float.h` C header file), `lower_bounds` will default to a vector of `-DBL_MAX` values, and `descriptors` will default to a vector of `'cdv_i'` strings, where `i` ranges from one to two:

```
variables,
    continuous_design = 2
```

In the following example, the syntax for a normal-lognormal distribution is shown. One normal and one lognormal uncertain variable are completely specified by their means and standard deviations. In addition, the dependence structure between the two variables is specified using the `uncertain_correlation_matrix`.

```
variables,
        normal_uncertain    =  1
          means             =  1.0
          std_deviations    =  1.0
          descriptors       =  'TF1n'
        lognormal_uncertain =  1
          means             =  2.0
          std_deviations    =  0.5
```

```
          descriptors       =  'TF2ln'
        uncertain_correlation_matrix =  1.0 0.2
                                        0.2 1.0
```

An example of the syntax for a state variables specification follows:

```
variables,
        continuous_state = 1
          initial_state       4.0
          lower_bounds        0.0
          upper_bounds        8.0
          descriptors        'CS1'
        discrete_state_range = 1
          initial_state       104
          lower_bounds        100
          upper_bounds        110
          descriptors        'DS1'
```

And in a more advanced example, a variables specification containing a set identifier, continuous and discrete design variables, normal and uniform uncertain variables, and continuous and discrete state variables is shown:

```
variables,
    id_variables = 'V1'
    continuous_design = 2
      initial_point    0.9    1.1
      upper_bounds     5.8    2.9
      lower_bounds     0.5   -2.9
      descriptors   'radius' 'location'
    discrete_design_range = 1
      initial_point    2
      upper_bounds     1
      lower_bounds     3
      descriptors   'material'
    normal_uncertain = 2
      means          =  248.89, 593.33
      std_deviations =   12.4,   29.7
      descriptors    =  'TF1n'   'TF2n'
    uniform_uncertain = 2
      lower_bounds =  199.3,  474.63
      upper_bounds =  298.5,  712.
      descriptors  =  'TF1u'   'TF2u'
    continuous_state = 2
      initial_state = 1.e-4  1.e-6
      descriptors   = 'EPSIT1' 'EPSIT2'
    discrete_state_set_int = 1
      initial_state = 100
      set_values    = 100 212 375
      descriptors   = 'load_case'
```

5.4.1 id_variables

- Keywords Area
- variables
- id_variables

Give the block a name, in case of multiple variables blocks

Topics

This keyword is related to the topics:

- block_pointer

Specification

Alias: none
Argument(s): STRING

Description

The optional set identifier specification uses the keyword `id_variables` to input a unique string for use in identifying a particular variables set. A model can then identify the use of this variables set by specifying the same string in its `variables_pointer` specification.

If the `id_variables` specification is omitted, a particular variables set will be used by a model only if that model omits specifying a `variables_pointer` and if the variables set was the last set parsed (or is the only set parsed). In common practice, if only one variables set exists, then `id_variables` can be safely omitted from the variables specification and `variables_pointer` can be omitted from the model specification(s), since there is no potential for ambiguity in this case.

Examples

For example, a model whose specification contains `variables_pointer = 'V1'` will use a variables specification containing the set identifier `id_variables = 'V1'`.

See Also

These keywords may also be of interest:

- variables_pointer

5.4.2 active

- Keywords Area
- variables
- active

Modify how the iterator treats different types of variables

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

			all	Option for the `active` keyword
	Required*(Choose One)*	**Group 1**	design	Option for the `active` keyword
			uncertain	Option for the `active` keyword
			aleatory	Option for the `active` keyword
			epistemic	Option for the `active` keyword
			state	Option for the `active` keyword

Description

There are certain situations where the user may want to explicitly control the subset of variables that is considered active for a certain Dakota method. This is done by specifying the keyword `active` in the variables specification block, followed by one of the following: `all`, `design`, `uncertain`, `aleatory`, `epistemic`, or `state`.

Specifying one of these subsets of variables will allow the Dakota method to operate on the specified variable types and override the default active subset.

If the user does not specify any explicit override of the active view of the variables, Dakota first considers the response function specification.

- If the user specifies objective functions or calibration terms in the response specification block, then we can infer that the active variables should be the design variables (since design variables are used within optimization and least squares methods).
- If the user instead specifies the generic response type of `response_functions`, then Dakota cannot infer the active variable subset from the response specification and will instead infer it from the method selection.
    1. If the method is a parameter study, or any of the methods available under dace, psuade, or fsu methods, the active view is set to all variables.
    2. For uncertainty quantification methods, if the method is sampling, then the view is set to aleatory if only aleatory variables are present, epistemic if only epistemic variables are present, or uncertain (covering both aleatory and epistemic) if both are present.
    3. If the uncertainty method involves interval estimation or evidence calculations, the view is set to epistemic.
    4. For other uncertainty quantification methods not mentioned in the previous sentences (e.g., reliability methods or stochastic expansion methods), the default view is set to aleatory.
    5. Finally, for verification studies using the Richardson extrapolation method, the active view is set to state.
    6. Note that in surrogate-based optimization, where the surrogate is built on points defined by the method defined by the `dace_method_pointer`, the sampling used to generate the points is performed only over the design variables as a default unless otherwise specified (e.g. state variables will not be sampled for surrogate construction).

As alluded to in the previous section, the iterative method selected for use in Dakota determines what subset, or view, of the variables data is active in the iteration. The general case of having a mixture of various different types of variables is supported within all of the Dakota methods even though certain methods will only modify

certain types of variables (e.g., optimizers and least squares methods only modify design variables, and uncertainty quantification methods typically only utilize uncertain variables). This implies that variables which are not under the direct control of a particular iterator will be mapped through the interface in an unmodified state. This allows for a variety of parameterizations within the model in addition to those which are being used by a particular iterator, which can provide the convenience of consolidating the control over various modeling parameters in a single file (the Dakota input file). An important related point is that the variable set that is active with a particular iterator is the same variable set for which derivatives are typically computed.

Examples

For example, the default behavior for a nondeterministic sampling method is to sample the uncertain variables. However, if the user specifed `active all` in the variables specification block, the sampling would be performed over all variables (e.g. design and state variables in addition to the uncertain variables). This may be desired in situations such as surrogate-based optimization under uncertainty, where a surrogate may be constructed to span both design and uncertain variables. This is an example where we expand the active subset beyond the default, but in other situations, we may wish to restrict from the default. An example of this would be performing design of experiments in the presence of multiple variable types (for which all types are active by default), but only wanting to sample over the design variables for purposes of constructing a surrogate model for optimization.

Theory

The optional status of the different variable type specifications allows the user to specify only those variables which are present (rather than explicitly specifying that the number of a particular type of variables is zero). However, at least one type of variables that are active for the iterator in use must have nonzero size or an input error message will result.

all

- Keywords Area
- variables
- active
- all

Option for the `active` keyword

Specification

Alias: none
Argument(s): none

Description

See the active keyword

design

- Keywords Area
- variables
- active
- design

Option for the `active` keyword

Specification

Alias: none
Argument(s): none

Description

See the active keyword

uncertain

- Keywords Area
- variables
- active
- uncertain

Option for the `active` keyword

Specification

Alias: none
Argument(s): none

Description

See the active keyword

aleatory

- Keywords Area
- variables
- active
- aleatory

Option for the `active` keyword

Specification

Alias: none
Argument(s): none

Description

See the active keyword

epistemic

- Keywords Area
- variables
- active
- epistemic

Option for the `active` keyword

Specification

Alias: none
Argument(s): none

Description

See the active keyword

state

- Keywords Area
- variables
- active
- state

Option for the `active` keyword

Specification

Alias: none
Argument(s): none

Description

See the active keyword

5.4.3 mixed

- Keywords Area
- variables
- mixed

Continuous and discrete variables are treated separately

Specification

Alias: none
Argument(s): none

Description

The variables domain specifies how the discrete variables are treated. If the user specifies `mixed` in the variable specification block, the continuous and discrete variables are treated separately. If the user specifies `relaxed` in the variable specification block, the discrete variables are relaxed and treated as continuous variables. This may be useful in optimization problems involving both continuous and discrete variables when a user would like to use an optimization method that is designed for continuous variable optimization. All Dakota methods have a default value of mixed for the domain type except for the branch-and-bound method which has a default domain type of relaxed. Note that the branch-and-bound method is under development at this time. Finally, note that the domain selection applies to all variable types: design, aleatory uncertain, epistemic uncertain, and state.

With respect to domain type, if the user does not specify an explicit override of `mixed` or `relaxed`, Dakota infers the domain type from the method. As mentioned above, all methods currently use a mixed domain as a default, except the branch-and-bound method which is under development.

See Also

These keywords may also be of interest:

- relaxed

5.4.4 relaxed

- Keywords Area
- variables
- relaxed

Treat discrete variables as continuous

Specification

Alias: none
Argument(s): none

Description

The variables domain specifies how the discrete variables are treated. If the user specifies `mixed` in the variable specification block, the continuous and discrete variables are treated separately. If the user specifies `relaxed` in the variable specification block, the discrete variables are relaxed and treated as continuous variables. This may be useful in optimization problems involving both continuous and discrete variables when a user would like to use an optimization method that is designed for continuous variable optimization. All Dakota methods have a default value of mixed for the domain type except for the branch-and-bound method which has a default domain type of relaxed. Note that the branch-and-bound method is under development at this time. Finally, note that the domain selection applies to all variable types: design, aleatory uncertain, epistemic uncertain, and state.

With respect to domain type, if the user does not specify an explicit override of `mixed` or `relaxed`, Dakota infers the domain type from the method. As mentioned above, all methods currently use a mixed domain as a default, except the branch-and-bound method which is under development.

See Also

These keywords may also be of interest:

- mixed

5.4.5 continuous_design

- Keywords Area
- variables
- continuous_design

Continuous design variables, each defined by a single interval

Topics

This keyword is related to the topics:

- continuous_variables
- design_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_point	Initial values
	Optional		lower_bounds	Specify minimum values

	Optional	upper_bounds	Specify maximium values
	Optional	scale_types	Specify scaling for the variables
	Optional	scales	Specify scaling for the variable.
	Optional	descriptors	Label the variables

Description

Continuous variables that are changed during the search for the optimal design.

initial_point

- Keywords Area
- variables
- continuous_design
- initial_point

Initial values

Specification

Alias: cdv_initial_point
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

lower_bounds

- Keywords Area
- variables
- continuous_design
- lower_bounds

Specify minimum values

Specification

Alias: cdv_lower_bounds
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- continuous_design
- upper_bounds

Specify maximium values

Specification

Alias: cdv_upper_bounds
Argument(s): REALLIST

Description

Specify maximium values

scale_types

- Keywords Area
- variables
- continuous_design
- scale_types

Specify scaling for the variables

Specification

Alias: cdv_scale_types
Argument(s): STRINGLIST

Description

For continuous variables, the `scale_types` specification includes strings specifying the scaling type for each component of the continuous design variables vector in methods that support scaling, when scaling is enabled.

Each entry in `scale_types` may be selected from `'none'`, `'value'`, `'auto'`, or `'log'`, to select no, characteristic value, automatic, or logarithmic scaling, respectively. If a single string is specified it will apply to all components of the continuous design variables vector. Each entry in `scales` may be a user-specified nonzero real characteristic value to be used in scaling each variable component. These values are ignored for scaling type `'none'`, required for `'value'`, and optional for `'auto'` and `'log'`. If a single real value is specified it will apply to all components of the continuous design variables vector.

Examples

Two continuous design variables, one scaled by a value, the other log scaled,

```
continuous_design = 2
  initial_point    -1.2      1.0
  lower_bounds     -2.0      0.001
  upper_bounds      2.0      2.0
  descriptors       'x1'     "x2"
  scale_types =     'value' 'log'
  scales =          4.0      0.1
```

scales

- Keywords Area
- variables
- continuous_design
- scales

Specify scaling for the variable.

Specification

Alias: cdv_scales
Argument(s): REALLIST

Description

For continuous variables, the `scale_types` specification includes strings specifying the scaling type for each component of the continuous design variables vector in methods that support scaling, when scaling is enabled. Each entry in `scale_types` may be selected from `'none'`, `'value'`, `'auto'`, or `'log'`, to select no, characteristic value, automatic, or logarithmic scaling, respectively. If a single string is specified it will apply to all components of the continuous design variables vector. Each entry in `scales` may be a user-specified nonzero real characteristic value to be used in scaling each variable component. These values are ignored for scaling type `'none'`, required for `'value'`, and optional for `'auto'` and `'log'`. If a single real value is specified it will apply to all components of the continuous design variables vector.

Examples

Two continuous design variables, both scaled by the characteristic value 4.0

```
continuous_design = 2
  initial_point    -1.2      1.0
  lower_bounds     -200      0.001
  upper_bounds      200      2.0
  descriptors       'x1'     "x2"
  scale_types = 'value' 'none'
  scales = 10.0
```

descriptors

- Keywords Area
- variables
- continuous_design
- descriptors

Label the variables

Specification

Alias: cdv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.6 discrete_design_range

- Keywords Area
- variables
- discrete_design_range

Discrete design variable - specified using a range of integers

Topics

This keyword is related to the topics:

- discrete_variables
- design_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_point	Initial values
	Optional		lower_bounds	Specify minimum values

	Optional	upper_bounds	Specify maximium values
	Optional	descriptors	Label the variables

Description

The details of how to specify this discrete variable are located on the discrete_variables page.

initial_point

- Keywords Area
- variables
- discrete_design_range
- initial_point

Initial values

Specification

Alias: ddv_initial_point
Argument(s): INTEGERLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

lower_bounds

- Keywords Area
- variables
- discrete_design_range
- lower_bounds

Specify minimum values

Specification

Alias: ddv_lower_bounds
Argument(s): INTEGERLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- discrete_design_range
- upper_bounds

Specify maximium values

Specification

Alias: ddv_upper_bounds
Argument(s): INTEGERLIST

Description

Specify maximium values

descriptors

- Keywords Area
- variables
- discrete_design_range
- descriptors

Label the variables

Specification

Alias: ddv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.7 discrete_design_set_integer

- Keywords Area
- variables
- discrete_design_set_integer

Discrete design variables - specified using lists of integers

Topics

This keyword is related to the topics:

- discrete_variables
- design_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_point	Initial values
	Optional		num_set_values	The size of each variables' set of permissible values
	Required		set_values	The permissible values for each discrete variable
	Optional		descriptors	Label the variables

Description

The details of how to specify this discrete variable are located on the discrete_variables page.

Examples

Four integer variables whose values will be selected during the search for an optimal design. $y1 \in [0, 1]$, $y2 \in [0, 1]$, $y3 \in [0, 5]$ and $y4 \in [10, 15, 20, 23]$.

```
discrete_design_set_integer 4
        num_set_values 2 2 2 4
        set_values 0 1 0 1 0 5 10 15 20 23
        descriptors 'y1' 'y2' 'y3' 'y4'
```

initial_point

- Keywords Area
- variables
- discrete_design_set_integer
- initial_point

Initial values

Specification

Alias: none

Argument(s): INTEGERLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

num_set_values

- Keywords Area
- variables
- discrete_design_set_integer
- num_set_values

The size of each variables' set of permissible values

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGERLIST

Description

For each variable, the number of permissible values must be specified.

set_values

- Keywords Area
- variables
- discrete_design_set_integer
- set_values

The permissible values for each discrete variable

Specification

Alias: none
Argument(s): INTEGERLIST

Description

This keyword allows specification of the permissible values for discrete set variables. All the variables of the type use this single keyword, so the specification can be confusing. See the description on the discrete_variables page.

descriptors

- Keywords Area
- variables
- discrete_design_set_integer
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.8 discrete_design_set_real

- Keywords Area
- variables
- discrete_design_set_real

Discrete design variables - specified using lists of reals

Topics

This keyword is related to the topics:

- discrete_variables
- design_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_point	Initial values
	Optional		num_set_values	The size of each variables' set of permissible values

	Required	set_values	The permissible values for each discrete variable
	Optional	descriptors	Label the variables

Description

The details of how to specify this discrete variable are located on the discrete_variables page.

Examples

Two continuous, restricted variables whose values will be selected during the search for an optimal design. $y1 \in [0.25, 1.25, 2.25, 3.25, 4.25]$, $y2 \in [0, 5]$

```
discrete_design_set_real 2
        num_set_values 5 2
        set_values 0.25 1.25 2.25 3.25 4.25 0 5
        descriptors 'y1' 'y2'
```

initial_point

- Keywords Area
- variables
- discrete_design_set_real
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

num_set_values

- Keywords Area
- variables
- discrete_design_set_real
- num_set_values

The size of each variables' set of permissible values

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGERLIST

Description

For each variable, the number of permissible values must be specified.

set_values

- Keywords Area
- variables
- discrete_design_set_real
- set_values

The permissible values for each discrete variable

Specification

Alias: none
Argument(s): REALLIST

Description

This keyword allows specification of the permissible values for discrete set variables. All the variables of the type use this single keyword, so the specification can be confusing. See the description on the discrete_variables page.

descriptors

- Keywords Area
- variables
- discrete_design_set_real
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.9 normal_uncertain

- Keywords Area
- variables
- normal_uncertain

Aleatory uncertain variable - normal (Gaussian)

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		means	First parameter of the distribution
	Required		std_deviations	Second parameter of the distribution
	Optional		lower_bounds	Specify minimum values
	Optional		upper_bounds	Specify maximium values
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

Within the normal uncertain optional group specification, the number of normal uncertain variables, the means, and standard deviations are required specifications, and the distribution lower and upper bounds and variable descriptors are optional specifications. The normal distribution is widely used to model uncertain variables such as population characteristics. It is also used to model the mean of a sample: as the sample size becomes very large, the Central Limit Theorem states that the distribution of the mean becomes approximately normal, regardless of the distribution of the original variables.

The density function for the normal distribution is:

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma_N} e^{-\frac{1}{2}\left(\frac{x-\mu_N}{\sigma_N}\right)^2}$$

where μ_N and σ_N are the mean and standard deviation of the normal distribution, respectively.

Note that if you specify bounds for a normal distribution, the sampling occurs from the underlying distribution with the given mean and standard deviation, but samples are not taken outside the bounds (see "bounded normal" distribution type in [[90] "Wyss and Jorgensen, 1998"]). This can result in the mean and the standard deviation of the sample data being different from the mean and standard deviation of the underlying distribution. For example, if you are sampling from a normal distribution with a mean of 5 and a standard deviation of 3, but you specify bounds of 1 and 7, the resulting mean of the samples will be around 4.3 and the resulting standard deviation will be around 1.6. This is because you have bounded the original distribution significantly, and asymetrically, since 7 is closer to the original mean than 1.

Theory

When used with design of experiments and multidimensional parameter studies, distribution bounds are inferred. These bounds are [$\mu - 3\sigma$, $\mu + 3\sigma$]

For vector and centered parameter studies, an inferred initial starting point is needed for the uncertain variables. These variables are initialized to their means for these studies.

means

- Keywords Area
- variables
- normal_uncertain
- means

First parameter of the distribution

Specification

Alias: nuv_means
Argument(s): REALLIST

Description

Means

std_deviations

- Keywords Area
- variables
- normal_uncertain
- std_deviations

Second parameter of the distribution

Specification

Alias: nuv_std_deviations
Argument(s): REALLIST

Description

Standard deviation

lower_bounds

- Keywords Area
- variables
- normal_uncertain
- lower_bounds

Specify minimum values

Specification

Alias: nuv_lower_bounds
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- normal_uncertain
- upper_bounds

Specify maximium values

Specification

Alias: nuv_upper_bounds
Argument(s): REALLIST

Description

Specify maximium values

initial_point

- Keywords Area
- variables
- normal_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- normal_uncertain
- descriptors

Label the variables

Specification

Alias: nuv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.10 lognormal_uncertain

- Keywords Area
- variables
- lognormal_uncertain

Aleatory uncertain variable - lognormal

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	lambdas	First parameter of the lognormal distribution (option 3)
			means	First parameter of the lognormal distribution (options 1 & 2)
	Optional		lower_bounds	Specify minimum values
	Optional		upper_bounds	Specify maximium values
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

If the logarithm of an uncertain variable X has a normal distribution, that is $\log X \sim N(\mu, \sigma)$, then X is distributed with a lognormal distribution. The lognormal is often used to model:

1. time to perform some task
2. variables which are the product of a large number of other quantities, by the Central Limit Theorem
3. quantities which cannot have negative values.

Within the lognormal uncertain optional group specification, the number of lognormal uncertain variables, the means, and either standard deviations or error factors must be specified, and the distribution lower and upper bounds and variable descriptors are optional specifications. These distribution bounds can be used to truncate the tails of lognormal distributions, which as for bounded normal, can result in the mean and the standard deviation of the sample data being different from the mean and standard deviation of the underlying distribution (see "bounded lognormal" and "bounded lognormal-n" distribution types in [[90] "Wyss and Jorgensen, 1998"]).

For the lognormal variables, one may specify either the mean μ and standard deviation σ of the actual lognormal distribution (option 1), the mean μ and error factor ϵ of the actual lognormal distribution (option 2), or the mean λ ("lambda") and standard deviation ζ ("zeta") of the underlying normal distribution (option 3).

The conversion equations from lognormal mean μ and either lognormal error factor ϵ or lognormal standard deviation σ to the mean λ and standard deviation ζ of the underlying normal distribution are as follows:

$$\zeta = \frac{ln(\epsilon)}{1.645}$$

$$\zeta^2 = ln(\frac{\sigma^2}{\mu^2} + 1)$$

$$\lambda = ln(\mu) - \frac{\zeta^2}{2}$$

Conversions from λ and ζ back to μ and ϵ or σ are as follows:

$$\mu = e^{\lambda + \frac{\zeta^2}{2}}$$

$$\sigma^2 = e^{2\lambda+\zeta^2}(e^{\zeta^2} - 1)$$
$$\epsilon = e^{1.645\zeta}$$

The density function for the lognormal distribution is:

$$f(x) = \frac{1}{\sqrt{2\pi}\zeta x} e^{-\frac{1}{2}\left(\frac{lnx-\lambda}{\zeta}\right)^2}$$

Theory

When used with design of experiments and multidimensional parameter studies, distribution bounds are inferred. These bounds are [0, $\mu + 3\sigma$].

For vector and centered parameter studies, an inferred initial starting point is needed for the uncertain variables. These variables are initialized to their means for these studies.

lambdas

- Keywords Area
- variables
- lognormal_uncertain
- lambdas

First parameter of the lognormal distribution (option 3)

Specification

Alias: lnuv_lambdas

Argument(s): REALLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		zetas	Second parameter of the lognormal distribution (option 3)

Description

For the lognormal variables, one may specify the mean λ ("lambda") and standard deviation ζ ("zeta") of the underlying normal distribution.

zetas

- Keywords Area
- variables
- lognormal_uncertain
- lambdas
- zetas

Second parameter of the lognormal distribution (option 3)

Specification

Alias: lnuv_zetas
Argument(s): REALLIST

Description

For the lognormal variables, one may specify the mean λ ("lambda") and standard deviation ζ ("zeta") of the underlying normal distribution.

means

- Keywords Area
- variables
- lognormal_uncertain
- means

First parameter of the lognormal distribution (options 1 & 2)

Specification

Alias: lnuv_means
Argument(s): REALLIST

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** std_deviations	**Dakota Keyword Description** Second parameter of the lognormal distribution (option 1)
			error_factors	Second parameter of the lognormal distribution (option 2)

Description

For the lognormal variables, one may specify either the mean μ and standard deviation σ of the actual lognormal distribution, the mean μ and error factor ϵ of the actual lognormal distribution.

This corresponds to the mean of the lognormal random variable

std_deviations

- Keywords Area
- variables
- lognormal_uncertain
- means
- std_deviations

Second parameter of the lognormal distribution (option 1)

Specification

Alias: lnuv_std_deviations
Argument(s): REALLIST

Description

For the lognormal variables, one may specify either the mean μ and standard deviation σ of the actual lognormal distribution.

This corresponds to the standard deviation of the lognormal random variable.

error_factors

- Keywords Area
- variables
- lognormal_uncertain
- means
- error_factors

Second parameter of the lognormal distribution (option 2)

Specification

Alias: lnuv_error_factors
Argument(s): REALLIST

Description

For the lognormal variables, one may specify the mean μ and error factor ϵ of the actual lognormal distribution.

This specifies the error function of the lognormal random variable.

lower_bounds

- Keywords Area
- variables
- lognormal_uncertain
- lower_bounds

Specify minimum values

Specification

Alias: lnuv_lower_bounds
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- lognormal_uncertain
- upper_bounds

Specify maximium values

Specification

Alias: lnuv_upper_bounds
Argument(s): REALLIST

Description

Specify maximium values

initial_point

- Keywords Area
- variables
- lognormal_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- lognormal_uncertain
- descriptors

Label the variables

Specification

Alias: lnuv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.11 uniform_uncertain

- Keywords Area
- variables
- uniform_uncertain

Aleatory uncertain variable - uniform

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		lower_bounds	Specify minimum values

	Required	upper_bounds	Specify maximium values
	Optional	initial_point	Initial values
	Optional	descriptors	Label the variables

Description

Within the uniform uncertain optional group specification, the number of uniform uncertain variables and the distribution lower and upper bounds are required specifications, and variable descriptors is an optional specification. The uniform distribution has the density function:

$$f(x) = \frac{1}{U_U - L_U}$$

where U_U and L_U are the upper and lower bounds of the uniform distribution, respectively. The mean of the uniform distribution is $\frac{U_U+L_U}{2}$ and the variance is $\frac{(U_U-L_U)^2}{12}$.

Theory

Note that this distribution is a special case of the more general beta distribution.

lower_bounds

- Keywords Area
- variables
- uniform_uncertain
- lower_bounds

Specify minimum values

Specification

Alias: uuv_lower_bounds
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- uniform_uncertain
- upper_bounds

Specify maximium values

Specification

Alias: uuv_upper_bounds
Argument(s): REALLIST

Description

Specify maximium values

initial_point

- Keywords Area
- variables
- uniform_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- uniform_uncertain
- descriptors

Label the variables

Specification

Alias: uuv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.12 loguniform_uncertain

- Keywords Area
- variables
- loguniform_uncertain

Aleatory uncertain variable - loguniform

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		lower_bounds	Specify minimum values
	Required		upper_bounds	Specify maximium values
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

If the logarithm of an uncertain variable X has a uniform distribution, that is $\log X \sim U(L_{LU}, U_{LU})$, then X is distributed with a loguniform distribution. The distribution lower bound is L_{LU} and upper bound is L_{LU} The loguniform distribution has the density function:

$$f(x) = \frac{1}{x(lnU_{LU} - lnL_{LU})}$$

Theory

For vector and centered parameter studies, an inferred initial starting point is needed for the uncertain variables. These variables are initialized to their means for these studies.

lower_bounds

- Keywords Area
- variables
- loguniform_uncertain
- lower_bounds

Specify minimum values

Specification

Alias: luuv_lower_bounds
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- loguniform_uncertain
- upper_bounds

Specify maximium values

Specification

Alias: luuv_upper_bounds
Argument(s): REALLIST

Description

Specify maximium values

initial_point

- Keywords Area
- variables
- loguniform_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- loguniform_uncertain
- descriptors

Label the variables

Specification

Alias: luuv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.13 triangular_uncertain

- Keywords Area
- variables
- triangular_uncertain

Aleatory uncertain variable - triangular

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		modes	Distribution parameter

	Required	lower_bounds	Specify minimum values
	Required	upper_bounds	Specify maximium values
	Optional	initial_point	Initial values
	Optional	descriptors	Label the variables

Description

The triangular distribution is often used when one does not have much data or information, but does have an estimate of the most likely value and the lower and upper bounds. Within the triangular uncertain optional group specification, the number of triangular uncertain variables, the modes, and the distribution lower and upper bounds are required specifications, and variable descriptors is an optional specification.

The density function for the triangular distribution is:

$$f(x) = \frac{2(x - L_T)}{(U_T - L_T)(M_T - L_T)}$$

if $L_T \leq x \leq M_T$, and

$$f(x) = \frac{2(U_T - x)}{(U_T - L_T)(U_T - M_T)}$$

if $M_T \leq x \leq U_T$, and 0 elsewhere. In these equations, L_T is the lower bound, U_T is the upper bound, and M_T is the mode of the triangular distribution.

modes

- Keywords Area
- variables
- triangular_uncertain
- modes

Distribution parameter

Specification

Alias: tuv_modes

Argument(s): REALLIST

Description

Specify the modes

lower_bounds

- Keywords Area
- variables
- triangular_uncertain
- lower_bounds

Specify minimum values

Specification

Alias: tuv_lower_bounds
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- triangular_uncertain
- upper_bounds

Specify maximium values

Specification

Alias: tuv_upper_bounds
Argument(s): REALLIST

Description

Specify maximium values

initial_point

- Keywords Area
- variables
- triangular_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- triangular_uncertain
- descriptors

Label the variables

Specification

Alias: tuv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.14 exponential_uncertain

- Keywords Area
- variables
- exponential_uncertain

Aleatory uncertain variable - exponential

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		betas	Parameter of the exponential distribution

	Optional	initial_point	Initial values
	Optional	descriptors	Label the variables

Description

The exponential distribution is often used for modeling failure rates.

The density function for the exponential distribution is given by:

$$f(x) = \frac{1}{\beta} e^{\frac{-x}{\beta}}$$

where $\mu_E = \beta$ and $\sigma_E^2 = \beta^2$.

Note that this distribution is a special case of the more general gamma distribution.

Theory

When used with design of experiments and multidimensional parameter studies, distribution bounds are inferred. These bounds are $[0, \mu + 3\sigma]$.

For vector and centered parameter studies, an inferred initial starting point is needed for the uncertain variables. These variables are initialized to their means for these studies.

betas

- Keywords Area
- variables
- exponential_uncertain
- betas

Parameter of the exponential distribution

Specification

Alias: euv_betas

Argument(s): REALLIST

Description

Specifies the list of β parameters to define the distributions of the exponential random variables. Length must match the other parameters and the number of exponential random variables.

initial_point

- Keywords Area
- variables
- exponential_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- exponential_uncertain
- descriptors

Label the variables

Specification

Alias: euv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.15 beta_uncertain

- Keywords Area
- variables
- beta_uncertain

Aleatory uncertain variable - beta

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		alphas	First parameter of the beta distribution
	Required		betas	Second parameter of the beta distribution
	Required		lower_bounds	Specify minimum values
	Required		upper_bounds	Specify maximium values
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

Within the beta uncertain optional group specification, the number of beta uncertain variables, the alpha and beta parameters, and the distribution upper and lower bounds are required specifications, and the variable descriptors is an optional specification. The beta distribution can be helpful when the actual distribution of an uncertain variable is unknown, but the user has a good idea of the bounds, the mean, and the standard deviation of the uncertain variable. The density function for the beta distribution is

$$f(x) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \frac{(x-L_B)^{\alpha-1}(U_B-x)^{\beta-1}}{(U_B-L_B)^{\alpha+\beta-1}}$$

where $\Gamma(\alpha)$ is the gamma function and $B(\alpha,\beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$ is the beta function. To calculate the mean and standard deviation from the alpha, beta, upper bound, and lower bound parameters of the beta distribution, the following expressions may be used.

$$\mu_B = L_B + \frac{\alpha}{\alpha+\beta}(U_B - L_B)$$

$$\sigma_B^2 = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}(U_B-L_B)^2$$

Solving these for α and β gives:

$$\alpha = (\mu_B - L_B)\frac{(\mu_B-L_B)(U_B-\mu_B)-\sigma_B^2}{\sigma_B^2(U_B-L_B)}$$

$$\beta = (U_B - \mu_B)\frac{(\mu_B-L_B)(U_B-\mu_B)-\sigma_B^2}{\sigma_B^2(U_B-L_B)}$$

Note that the uniform distribution is a special case of this distribution for parameters $\alpha = \beta = 1$.

Theory

For vector and centered parameter studies, an inferred initial starting point is needed for the uncertain variables. These variables are initialized to their means for these studies.

alphas

- Keywords Area
- variables
- beta_uncertain
- alphas

First parameter of the beta distribution

Specification

Alias: buv_alphas
Argument(s): REALLIST

Description

Specifies the list of α parameters to define the distributions of the beta random variables. Length must match the other parameters and the number of beta random variables.

betas

- Keywords Area
- variables
- beta_uncertain
- betas

Second parameter of the beta distribution

Specification

Alias: buv_betas
Argument(s): REALLIST

Description

Specifies the list of β parameters to define the distributions of the beta random variables. Length must match the other parameters and the number of beta random variables.

lower_bounds

- Keywords Area
- variables
- beta_uncertain
- lower_bounds

Specify minimum values

Specification

Alias: buv_lower_bounds
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- beta_uncertain
- upper_bounds

Specify maximium values

Specification

Alias: buv_upper_bounds
Argument(s): REALLIST

Description

Specify maximium values

initial_point

- Keywords Area
- variables
- beta_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- beta_uncertain
- descriptors

Label the variables

Specification

Alias: buv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.16 gamma_uncertain

- Keywords Area
- variables
- gamma_uncertain

Aleatory uncertain variable - gamma

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		alphas	First parameter of the gamma distribution

	Required	betas	Second parameter of the gamma distribution
	Optional	initial_point	Initial values
	Optional	descriptors	Label the variables

Description

The gamma distribution is sometimes used to model time to complete a task, such as a repair or service task. It is a very flexible distribution with its shape governed by alpha and beta.

The density function for the gamma distribution is given by:

$$f(x) = \frac{x^{\alpha-1} e^{\frac{-x}{\beta}}}{\beta^{\alpha}\Gamma(\alpha)}$$

where $\mu_{GA} = \alpha\beta$ and $\sigma^2_{GA} = \alpha\beta^2$. Note that the exponential distribution is a special case of this distribution for parameter $\alpha = 1$.

Theory

When used with design of experiments and multidimensional parameter studies, distribution bounds are inferred. These bounds are [0, $\mu + 3\sigma$].

For vector and centered parameter studies, an inferred initial starting point is needed for the uncertain variables. These variables are initialized to their means for these studies.

alphas

- Keywords Area
- variables
- gamma_uncertain
- alphas

First parameter of the gamma distribution

Specification

Alias: gauv_alphas

Argument(s): REALLIST

Description

Specifies the list of α parameters to define the distributions of the gamma random variables. Length must match the other parameters and the number of gamma random variables.

betas

- Keywords Area
- variables
- gamma_uncertain
- betas

Second parameter of the gamma distribution

Specification

Alias: gauv_betas
Argument(s): REALLIST

Description

Specifies the list of β parameters to define the distributions of the gamma random variables. Length must match the other parameters and the number of gamma random variables.

initial_point

- Keywords Area
- variables
- gamma_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- gamma_uncertain
- descriptors

Label the variables

Specification

Alias: gauv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.17 gumbel_uncertain

- Keywords Area
- variables
- gumbel_uncertain

Aleatory uncertain variable - gumbel

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		alphas	First parameter of the gumbel distribution
	Required		betas	Second parameter of the gumbel distribution
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

The Gumbel distribution is also referred to as the Type I Largest Extreme Value distribution. The distribution of maxima in sample sets from a population with a normal distribution will asymptotically converge to this distribution. It is commonly used to model demand variables such as wind loads and flood levels.

The density function for the Gumbel distribution is given by:

$$f(x) = \alpha e^{-\alpha(x-\beta)} exp(-e^{-\alpha(x-\beta)})$$

where $\mu_{GU} = \beta + \frac{0.5772}{\alpha}$ and $\sigma_{GU} = \frac{\pi}{\sqrt{6}\alpha}$.

Theory

When used with design of experiments and multidimensional parameter studies, distribution bounds are inferred. These bounds are [$\mu - 3\sigma$, $\mu + 3\sigma$]

For vector and centered parameter studies, an inferred initial starting point is needed for the uncertain variables. These variables are initialized to their means for these studies.

alphas

- Keywords Area
- variables
- gumbel_uncertain
- alphas

First parameter of the gumbel distribution

Specification

Alias: guuv_alphas
Argument(s): REALLIST

Description

Specifies the list of β parameters to define the distributions of the gumbel random variables. Length must match the other parameters and the number of gumbel random variables.

betas

- Keywords Area
- variables
- gumbel_uncertain
- betas

Second parameter of the gumbel distribution

Specification

Alias: guuv_betas
Argument(s): REALLIST

Description

Specifies the list of β parameters to define the distributions of the gumbel random variables. Length must match the other parameters and the number of gumbel random variables.

initial_point

- Keywords Area
- variables
- gumbel_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- gumbel_uncertain
- descriptors

Label the variables

Specification

Alias: guuv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.18 frechet_uncertain

- Keywords Area
- variables
- frechet_uncertain

Aleatory uncertain variable - Frechet

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		alphas	First parameter of the Frechet distribution
	Required		betas	Second parameter of the Frechet distribution
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

The Frechet distribution is also referred to as the Type II Largest Extreme Value distribution. The distribution of maxima in sample sets from a population with a lognormal distribution will asymptotically converge to this distribution. It is commonly used to model non-negative demand variables.

The density function for the frechet distribution is:

$$f(x) = \frac{\alpha}{\beta}(\frac{\beta}{x})^{\alpha+1}e^{-(\frac{\beta}{x})^{\alpha}}$$

where $\mu_F = \beta\Gamma(1-\frac{1}{\alpha})$ and $\sigma_F^2 = \beta^2[\Gamma(1-\frac{2}{\alpha}) - \Gamma^2(1-\frac{1}{\alpha})]$

Theory

When used with design of experiments and multidimensional parameter studies, distribution bounds are inferred. These bounds are [0, $\mu + 3\sigma$].

For vector and centered parameter studies, an inferred initial starting point is needed for the uncertain variables. These variables are initialized to their means for these studies.

alphas

- Keywords Area
- variables
- frechet_uncertain
- alphas

First parameter of the Frechet distribution

Specification

Alias: fuv_alphas
Argument(s): REALLIST

Description

Specifies the list of α parameters to define the distributions of the Frechet random variables. Length must match the other parameters and the number of Frechet random variables.

betas

- Keywords Area
- variables
- frechet_uncertain
- betas

Second parameter of the Frechet distribution

Specification

Alias: fuv_betas
Argument(s): REALLIST

Description

Specifies the list of β parameters to define the distributions of the Frechet random variables. Length must match the other parameters and the number of Frechet random variables.

initial_point

- Keywords Area
- variables
- frechet_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- frechet_uncertain
- descriptors

Label the variables

Specification

Alias: fuv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.19 weibull_uncertain

- Keywords Area
- variables
- weibull_uncertain

Aleatory uncertain variable - Weibull

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		alphas	First parameter of the Weibull distribution

	Required	betas	Second parameter of the Weibull distribution
	Optional	initial_point	Initial values
	Optional	descriptors	Label the variables

Description

The Weibull distribution is also referred to as the Type III Smallest Extreme Value distribution. The Weibull distribution is commonly used in reliability studies to predict the lifetime of a device. It is also used to model capacity variables such as material strength.

The density function for the Weibull distribution is given by:

$$f(x) = \frac{\alpha}{\beta}\left(\frac{x}{\beta}\right)^{\alpha-1} e^{-\left(\frac{x}{\beta}\right)^{\alpha}}$$

where $\mu_W = \beta\Gamma(1+\frac{1}{\alpha})$ and $\sigma_W = \sqrt{\frac{\Gamma(1+\frac{2}{\alpha})}{\Gamma^2(1+\frac{1}{\alpha})} - 1}\mu_W$

alphas

- Keywords Area
- variables
- weibull_uncertain
- alphas

First parameter of the Weibull distribution

Specification

Alias: wuv_alphas

Argument(s): REALLIST

Description

Specifies the list of α parameters to define the distributions of the Weibull random variables.

Length must match the other parameters and the number of Weibull random variables.

betas

- Keywords Area
- variables
- weibull_uncertain
- betas

Second parameter of the Weibull distribution

Specification

Alias: wuv_betas
Argument(s): REALLIST

Description

Specifies the list of β parameters to define the distributions of the Weibull random variables. Length must match the other parameters and the number of Weibull random variables.

initial_point

- Keywords Area
- variables
- weibull_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- weibull_uncertain
- descriptors

Label the variables

Specification

Alias: wuv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.20 histogram_bin_uncertain

- Keywords Area
- variables
- histogram_bin_uncertain

Aleatory uncertain variable - an arbitrary continuous random variable !!!!!!!!!!!!!!!clarify difference between ordinatea and count

Topics

This keyword is related to the topics:

- continuous_variables
- aleatory_uncertain_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_pairs	???
	Required		abscissas	Distribution parameter
	Required*(Choose One)*	**Group 1**	ordinates	??????????
			counts	???
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

Histogram uncertain variables are typically used to model a set of empirical data. A bin histogram is a continuous aleatory distribution that allows the user to specify bins of non-zero width where the uncertain variable may lie along with the relative frequencies that are associated with each bin. When using a histogram bin variable, one must define at least one bin (with two bounding value pairs). One can characterize the bin with non-zero widths (abscissas and counts) or with individual points (abscissas and ordinates).

The `abscissas` specification define abscissa values (”x” coordinates) for the probability density function of each histogram variable. When paired with `counts`, the specifications provide sets of (x,c) pairs for each histogram variable where `c` defines a count (i.e., a frequency or relative probability) associated with a bin. If using bins of unequal width and specification of probability densities is more natural, then the `counts` specification can be replaced with a `ordinates` specification (”y” coordinates) in order to support interpretation of the input as (x,y) pairs defining the profile of a ”skyline” probability density function. Conversion between the two specifications is straightforward: a count/frequency is a cumulative probability quantity defined from the product of the ordinate density value and the `x` bin width. Thus, in the cases of bins of equal width, ordinate and count specifications are equivalent. In addition, ordinates and counts may be relative values; it is not necessary to scale them as all user inputs will be normalized.

To fully specify a bin-based histogram with `n` bins (potentially of unequal width), `n+1` (x,c) or (x,y) pairs must be specified with the following features:

- `x` is the parameter value for the left boundary of a histogram bin and `c` is the corresponding count for that bin. Alternatively, `y` defines the ordinate density value for this bin within a skyline probability density function. The right boundary of the bin is defined by the left boundary of the next pair.
- the final pair specifies the right end of the last bin and must have a `c` or `y` value of zero.
- the `x` values must be strictly increasing.
- all `c` or `y` values must be positive, except for the last which must be zero.
- a minimum of two pairs must be specified for each bin-based histogram.

The number of pairs specifications provide for the proper association of multiple sets of (x,c) or (x,y) pairs with individual histogram variables. For example, in the following specification

```
histogram_bin_uncertain = 2
 num_pairs = 3 4
 abscissas = 5 8 10 .1 .2 .3 .4
 counts = 17 21 0 12 24 12 0
```

`num_pairs` associates the first 3 (x,c) pairs from `abscissas` and `counts` ((5,17),(8,21),(10,0)) with one bin-based histogram variable, where one bin is defined between 5 and 8 with a count of 17 and another bin is defined between 8 and 10 with a count of 21. The following set of 4 (x,c) pairs ((.1,12),(.2,24),(.3,12),(.4,0)) defines a second bin-based histogram variable containing three equal-width bins with counts 12, 24, and 12 (middle bin is twice as probable as the other two).

num_pairs

- Keywords Area
- variables
- histogram_bin_uncertain
- num_pairs

???

Topics

This keyword is related to the topics:

- problem

Specification

Alias: huv_num_bin_pairs
Argument(s): INTEGERLIST

Description

Within the histogram bin uncertain optional group specification, the number of histogram bin uncertain variables is a required specification, the number of pairs is an optional key for apportionment of abscissas/ordinates/counts, specification of abscissas and either ordinates or counts is required, and the variable descriptors is an optional specification. When using a histogram bin variable, one must define at least one bin (with two bounding value pairs).

The `abscissas` specification define abscissa values (”x” coordinates) for the probability density function of each histogram variable. When paired with `counts`, the specifications provide sets of (x,c) pairs for each histogram variable where `c` defines a count (i.e., a frequency or relative probability) associated with a bin. If using bins of unequal width and specification of probability densities is more natural, then the `counts` specification can be replaced with a `ordinates` specification (”y” coordinates) in order to support interpretation of the input as (x,y) pairs defining the profile of a ”skyline” probability density function. Conversion between the two specifications is straightforward: a count/frequency is a cumulative probability quantity defined from the product of the ordinate density value and the `x` bin width. Thus, in the cases of bins of equal width, ordinate and count specifications are equivalent. In addition, ordinates and counts may be relative values; it is not necessary to scale them as all user inputs will be normalized.

abscissas

- Keywords Area
- variables
- histogram_bin_uncertain
- abscissas

Distribution parameter

Specification

Alias: huv_bin_abscissas
Argument(s): REALLIST

Description

Within the histogram bin uncertain optional group specification, the number of histogram bin uncertain variables is a required specification, the number of pairs is an optional key for apportionment of abscissas/ordinates/counts, specification of abscissas and either ordinates or counts is required, and the variable descriptors is an optional specification. When using a histogram bin variable, one must define at least one bin (with two bounding value pairs).

The `abscissas` specification define abscissa values (”x” coordinates) for the probability density function of each histogram variable. When paired with `counts`, the specifications provide sets of (x,c) pairs for each histogram variable where `c` defines a count (i.e., a frequency or relative probability) associated with a bin. If using bins of unequal width and specification of probability densities is more natural, then the `counts` specification can be replaced with a `ordinates` specification (”y” coordinates) in order to support interpretation of the input as (x,y) pairs defining the profile of a ”skyline” probability density function. Conversion between the two specifications is straightforward: a count/frequency is a cumulative probability quantity defined from the product of the ordinate density value and the `x` bin width. Thus, in the cases of bins of equal width, ordinate and count specifications are equivalent. In addition, ordinates and counts may be relative values; it is not necessary to scale them as all user inputs will be normalized.

ordinates

- Keywords Area
- variables
- histogram_bin_uncertain
- ordinates

??????????

Topics

This keyword is related to the topics:

- problem

Specification

Alias: huv_bin_ordinates
Argument(s): REALLIST

Description

Within the histogram bin uncertain optional group specification, the number of histogram bin uncertain variables is a required specification, the number of pairs is an optional key for apportionment of abscissas/ordinates/counts, specification of abscissas and either ordinates or counts is required, and the variable descriptors is an optional specification. When using a histogram bin variable, one must define at least one bin (with two bounding value pairs).

The `abscissas` specification define abscissa values (”x” coordinates) for the probability density function of each histogram variable. When paired with `counts`, the specifications provide sets of (x,c) pairs for each histogram variable where `c` defines a count (i.e., a frequency or relative probability) associated with a bin. If using bins of unequal width and specification of probability densities is more natural, then the `counts` specification can be replaced with a `ordinates` specification (”y” coordinates) in order to support interpretation of the input as (x,y) pairs defining the profile of a ”skyline” probability density function. Conversion between the two specifications is straightforward: a count/frequency is a cumulative probability quantity defined from the product of the ordinate density value and the `x` bin width. Thus, in the cases of bins of equal width, ordinate and count specifications are equivalent. In addition, ordinates and counts may be relative values; it is not necessary to scale them as all user inputs will be normalized.

counts

- Keywords Area
- variables
- histogram_bin_uncertain
- counts

???

Topics

This keyword is related to the topics:

- problem

Specification

Alias: huv_bin_counts
Argument(s): REALLIST

Description

Within the histogram bin uncertain optional group specification, the number of histogram bin uncertain variables is a required specification, the number of pairs is an optional key for apportionment of abscissas/ordinates/counts, specification of abscissas and either ordinates or counts is required, and the variable descriptors is an optional specification. When using a histogram bin variable, one must define at least one bin (with two bounding value pairs).

The `abscissas` specification define abscissa values (”x” coordinates) for the probability density function of each histogram variable. When paired with `counts`, the specifications provide sets of (x,c) pairs for each histogram variable where `c` defines a count (i.e., a frequency or relative probability) associated with a bin. If using bins of unequal width and specification of probability densities is more natural, then the `counts` specification can be replaced with a `ordinates` specification (”y” coordinates) in order to support interpretation of the input as (x,y) pairs defining the profile of a ”skyline” probability density function. Conversion between the two specifications is straightforward: a count/frequency is a cumulative probability quantity defined from the product of the ordinate density value and the `x` bin width. Thus, in the cases of bins of equal width, ordinate and count specifications are equivalent. In addition, ordinates and counts may be relative values; it is not necessary to scale them as all user inputs will be normalized.

initial_point

- Keywords Area
- variables
- histogram_bin_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- histogram_bin_uncertain
- descriptors

Label the variables

Specification

Alias: huv_bin_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.21 poisson_uncertain

- Keywords Area
- variables
- poisson_uncertain

Aleatory uncertain discrete variable - Poisson

Topics

This keyword is related to the topics:

- discrete_variables
- aleatory_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		lambdas	The parameter for the Poisson distribution, the expected number of events in the time interval of interest
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

The Poisson distribution is used to predict the number of discrete events that happen in a single time interval. The random events occur uniformly and independently. The expected number of occurences in a single time interval is λ, which must be a positive real number. For example, if events occur on average 4 times per year and we are interested in the distribution of events over six months, λ would be 2. However, if we were interested in the distribution of events occuring over 5 years, λ would be 20.

The density function for the poisson distribution is given by:

$$f(x) = \frac{\lambda e^{-\lambda}}{x!}$$

where

- λ is the expected number of events occuring in a single time interval -x is the number of events that occur in this time period -f(x) is the probability that x events occur in this time period

Theory

When used with design of experiments and multidimensional parameter studies, distribution bounds are inferred. These bounds are [0, $\mu + 3\sigma$].

For vector and centered parameter studies, an inferred initial starting point is needed for the uncertain variables. These variables are initialized to their means for these studies.

lambdas

- Keywords Area
- variables
- poisson_uncertain
- lambdas

The parameter for the Poisson distribution, the expected number of events in the time interval of interest

Specification

Alias: none

Argument(s): REALLIST

Description

The density function for the poisson distribution is given by:

$$f(x) = \frac{\lambda e^{-\lambda}}{x!}$$

where λ is the frequency of events happening, and `x` is the number of events that occur.

initial_point

- Keywords Area
- variables
- poisson_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- poisson_uncertain
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.22 binomial_uncertain

- Keywords Area
- variables
- binomial_uncertain

Aleatory uncertain discrete variable - binomial

Topics

This keyword is related to the topics:

- discrete_variables
- aleatory_uncertain_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		probability_per_-trial	A distribution parameter for the binomial distribution
	Required		num_trials	A distribution parameter
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

The binomial distribution describes probabilities associated with a series of independent Bernoulli trials. A Bernoulli trial is an event with two mutually exclusive outcomes, such as 0 or 1, yes or no, success or fail. The probability of success remains the same (the trials are independent).

The density function for the binomial distribution is given by:

$$f(x) = \left(\begin{array}{c} n \\ x \end{array} \right) p^x (1-p)^{(n-x)}$$

where `p` is the probability of failure per trial, `n` is the number of trials and `x` is the number of successes.

Theory

The binomial distribution is typically used to predict the number of failures or defective items in a total of `n` independent tests or trials, where each trial has the probability `p` of failing or being defective.

probability_per_trial

- Keywords Area
- variables
- binomial_uncertain
- probability_per_trial

A distribution parameter for the binomial distribution

Specification

Alias: prob_per_trial
Argument(s): REALLIST

Description

The binomial distribution is typically used to predict the number of failures (or defective items or some type of event) in a total of `n` independent tests or trials, where each trial has the probability `p` of failing or being defective. Each particular test can be considered as a Bernoulli trial.

num_trials

- Keywords Area
- variables
- binomial_uncertain
- num_trials

A distribution parameter

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The binomial distribution is typically used to predict the number of failures (or defective items or some type of event) in a total of `n` independent tests or trials, where each trial has the probability `p` of failing or being defective. Each particular test can be considered as a Bernoulli trial.

initial_point

- Keywords Area
- variables
- binomial_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- binomial_uncertain
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.23 negative_binomial_uncertain

- Keywords Area
- variables
- negative_binomial_uncertain

Aleatory uncertain discrete variable - negative binomial

Topics

This keyword is related to the topics:

- discrete_variables
- aleatory_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		probability_per_-trial	A negative binomial distribution parameter
	Required		num_trials	A negative binomial distribution parameter
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

The density function for the negative binomial distribution is given by:

$$f(x) = \left(\begin{array}{c} n+x-1 \\ x \end{array} \right) p^n (1-p)^x$$

where

- p is the probability of success per trial
- n is the number of successful trials
- X is the number of failures

Theory

The negative binomial distribution is typically used to predict the number of failures observed when repeating a test until a total of `n` successes have occurred, where each test has a probability `p` of success.

probability_per_trial

- Keywords Area
- variables
- negative_binomial_uncertain
- probability_per_trial

A negative binomial distribution parameter

Specification

Alias: prob_per_trial
Argument(s): REALLIST

Description

The negative binomial distribution is typically used to predict the number of failures observed when repeating a test until a total of `n` successes have occurred, where each test has a probability `p` of success.

The density function for the negative binomial distribution is given by:

$$f(x) = \binom{n+x-1}{x} p^n (1-p)^x$$

where

- p is the probability of success per trial
- n is the number of successful trials
- X is the number of failures

num_trials

- Keywords Area
- variables
- negative_binomial_uncertain
- num_trials

A negative binomial distribution parameter

Specification

Alias: none

Argument(s): INTEGERLIST

Description

The negative binomial distribution is typically used to predict the number of failures observed when repeating a test until a total of `n` successes have occurred, where each test has a probability `p` of success.

The density function for the negative binomial distribution is given by:

$$f(x) = \binom{n+x-1}{x} p^n (1-p)^x$$

where

- p is the probability of success per trial
- n is the number of successful trials
- X is the number of failures

initial_point

- Keywords Area
- variables
- negative_binomial_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- negative_binomial_uncertain
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.24 geometric_uncertain

- Keywords Area
- variables
- geometric_uncertain

Aleatory uncertain discrete variable - geometric

Topics

This keyword is related to the topics:

- discrete_variables
- aleatory_uncertain_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		probability_per_-trial	Geometric distribution parameter
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

The geometric distribution represents the number of successful trials that might occur before a failure is observed.

The density function for the geometric distribution is given by:

$$f(x) = p(1-p)^x$$

where p is the probability of failure per trial.

probability_per_trial

- Keywords Area
- variables
- geometric_uncertain
- probability_per_trial

Geometric distribution parameter

Specification

Alias: prob_per_trial

Argument(s): REALLIST

Description

The geometric distribution represents the number of successful trials that occur before a failure is observed.

The density function for the geometric distribution is given by:

$$f(x) = p(1-p)^x$$

where p is the probability of failure per trial and x is the number of successful trials.

initial_point

- Keywords Area
- variables
- geometric_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- geometric_uncertain
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.25 hypergeometric_uncertain

- Keywords Area
- variables
- hypergeometric_uncertain

Aleatory uncertain discrete variable - hypergeometric

Topics

This keyword is related to the topics:

- discrete_variables
- aleatory_uncertain_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		total_population	Parameter for the hypergeometric probability distribution
	Required		selected_-population	Distribution parameter for the hypergeometric distribution
	Required		num_drawn	Distribution parameter for the hypergeometric distribution
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

The hypergeometric probability density is used when sampling without replacement from a total population of elements where

- The resulting element of each sample can be separated into one of two non-overlapping sets
- The probability of success changes with each sample.

The density function for the hypergeometric distribution is given by:

$$f(x) = \frac{\begin{pmatrix} m \\ x \end{pmatrix} \begin{pmatrix} N-m \\ n-x \end{pmatrix}}{\begin{pmatrix} N \\ n \end{pmatrix}}$$

where:

- N is the total population
- m is the number of items in the selected population (e.g. the number of white balls in the full urn of N items)
- n is the size of the sample drawn (e.g. number of balls drawn)
- x is the number of sucess (e.g. drawing a white ball)
- $\frac{a}{b}$ is a binomial coeffecient

Theory

The hypergeometric is often described using an urn model. For example, say we have a total population containing N balls, and we know that m of the balls are white and the remaining balls are green. If we draw n balls from the urn without replacement, the hypergeometric distribution describes the probability of drawing x white balls.

total_population

- Keywords Area
- variables
- hypergeometric_uncertain
- total_population

Parameter for the hypergeometric probability distribution

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The density function for the hypergeometric distribution is given by:

$$f(x) = \frac{\begin{pmatrix} m \\ x \end{pmatrix} \begin{pmatrix} N-m \\ n-x \end{pmatrix}}{\begin{pmatrix} N \\ n \end{pmatrix}}$$

where

- N is the total population (e.g. the total number of balls in the urn)
- m is the number of items in the selected population (e.g. the number of white balls in the full urn of N items)
- n is the size of the sample (e.g. number of balls drawn)
- x is the number of sucess (e.g. drawing a white ball)
- $\frac{a}{b}$ is a binomial coeffecient

selected_population

- Keywords Area
- variables
- hypergeometric_uncertain
- selected_population

Distribution parameter for the hypergeometric distribution

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The density function for the hypergeometric distribution is given by:

$$f(x) = \frac{\begin{pmatrix} m \\ x \end{pmatrix} \begin{pmatrix} N-m \\ n-x \end{pmatrix}}{\begin{pmatrix} N \\ n \end{pmatrix}}$$

where

- N is the total population (e.g. the total number of balls in the urn)
- m is the number of items in the selected population (e.g. the number of white balls in the full urn of N items)
- n is the size of the sample (e.g. number of balls drawn)
- x is the number of sucess (e.g. drawing a white ball)
- $\frac{a}{b}$ is a binomial coeffecient

num_drawn

- Keywords Area
- variables
- hypergeometric_uncertain
- num_drawn

Distribution parameter for the hypergeometric distribution

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The density function for the hypergeometric distribution is given by:

$$f(x) = \frac{\begin{pmatrix} m \\ x \end{pmatrix} \begin{pmatrix} N-m \\ n-x \end{pmatrix}}{\begin{pmatrix} N \\ n \end{pmatrix}}$$

where

- N is the total population (e.g. the total number of balls in the urn)
- m is the number of items in the selected population (e.g. the number of white balls in the full urn of N items)
- n is the size of the sample (e.g. number of balls drawn)
- x is the number of sucess (e.g. drawing a white ball)
- $\frac{a}{b}$ is a binomial coeffecient

initial_point

- Keywords Area
- variables
- hypergeometric_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- hypergeometric_uncertain
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.26 histogram_point_uncertain

- Keywords Area
- variables
- histogram_point_uncertain

Aleatory uncertain discrete variable - arbitrary discrete random variable

Topics

This keyword is related to the topics:

- discrete_variables
- aleatory_uncertain_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_pairs	Distribution parameter
	Required		abscissas	Distribution parameter for discrete point histograms
	Required		counts	Distribution parameter for the histogram point distribution
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

Histogram uncertain variables are typically used to model a set of empirical data. A histogram can be viewed as a probability distribution function (PDF). A point histogram is a discrete aleatory distribution that allows the user to specify a set of real-valued points and associated frequency values or counts.

When using a histogram point variable, one must define at least one set of abscissa/count pairs. The `abscissas` define abscissa values ("x" coordinates) for the PDF of each histogram variable. These are paired with `counts` which provide sets of (x,c) pairs for each histogram variable where `c` defines a count (i.e., a frequency) associated with a point.

To fully specify a point-based histogram with `n` points, `n` (x,c) pairs must be specified with the following features:

- `x` is the point value and `c` is the corresponding count for that value.
- the `x` values must be strictly increasing.

- all `c` values must be positive.
- a minimum of one pair must be specified for each point-based histogram.

The `num_pairs` specification provides for the proper association of multiple sets of (x,c) or (x,y) pairs with individual histogram variables. For example, in the following specification,

```
histogram_point_uncertain = 2
  num_pairs =   2           3
  abscissas = 3 4 100 200 300
  counts    = 1 1   1   2   1
```

`num_pairs` associates the (x,c) pairs ((3,1),(4,1)) with one point-based histogram variable (where the values 3 and 4 are equally probable) and associates the (x,c) pairs ((100,1),(200,2),(300,1)) with a second point-based histogram variable (where the value 200 is twice as probable as either 100 or 300).

num_pairs

- Keywords Area
- variables
- histogram_point_uncertain
- num_pairs

Distribution parameter

Specification

Alias: huv_num_point_pairs
Argument(s): INTEGERLIST

Description

Within the histogram point uncertain optional group specification, the number of histogram point uncertain variables is a required specification, the number of pairs is an optional key for apportionment of abscissas and counts, the sets of abscissas and counts are required, and the variable descriptors are optional. The `num_pairs` specification provides for the proper association of multiple sets of (x,c) or (x,y) pairs with individual histogram variables.

abscissas

- Keywords Area
- variables
- histogram_point_uncertain
- abscissas

Distribution parameter for discrete point histograms

Specification

Alias: huv_point_abscissas
Argument(s): REALLIST

Description

When using a histogram point variable, one must define at least one set of abscissa/count pairs. The `abscissas` define abscissa values ("x" coordinates) for the PDF of each histogram variable. These are paired with `counts` which provide sets of (x,c) pairs for each histogram variable where `c` defines a count (i.e., a frequency) associated with a point.

To fully specify a point-based histogram with `n` points, `n` (x,c) pairs must be specified with the following features:

- `x` is the point value and `c` is the corresponding count for that value.
- the `x` values must be strictly increasing.
- all `c` values must be positive.
- a minimum of one pair must be specified for each point-based histogram.

counts

- Keywords Area
- variables
- histogram_point_uncertain
- counts

Distribution parameter for the histogram point distribution

Specification

Alias: huv_point_counts
Argument(s): REALLIST

Description

When using a histogram point variable, one must define at least one set of abscissa/count pairs. The `abscissas` define abscissa values ("x" coordinates) for the PDF of each histogram variable. These are paired with `counts` which provide sets of (x,c) pairs for each histogram variable where `c` defines a count (i.e., a frequency) associated with a point.

To fully specify a point-based histogram with `n` points, `n` (x,c) pairs must be specified with the following features:

- `x` is the point value and `c` is the corresponding count for that value.
- the `x` values must be strictly increasing.
- all `c` values must be positive.
- a minimum of one pair must be specified for each point-based histogram.

initial_point

- Keywords Area
- variables
- histogram_point_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- histogram_point_uncertain
- descriptors

Label the variables

Specification

Alias: huv_point_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.27 uncertain_correlation_matrix

- Keywords Area
- variables
- uncertain_correlation_matrix

Specify correlations between aleatory uncertain variables

Specification

Alias: none
Argument(s): REALLIST

Description

Aleatory uncertain variables may have correlations specified through use of an `uncertain_correlation_matrix` specification. This specification is generalized in the sense that its specific meaning depends on the nondeterministic method in use.

When the method is a nondeterministic sampling method (i.e., sampling), then the correlation matrix specifies *rank correlations* [[52] "Iman and Conover, 1982"].

When the method is a reliability (i.e., `local_reliability` or `global_reliability`) or stochastic expansion (i.e., `polynomial_chaos` or `stoch_collocation`) method, then the correlation matrix specifies *correlation coefficients* (normalized covariance) [[42] "Haldar and Mahadevan, 2000"].

In either of these cases, specifying the identity matrix results in uncorrelated uncertain variables (the default). The matrix input should be symmetric and have all n^2 entries where n is the total number of aleatory uncertain variables.

Ordering of the aleatory uncertain variables is:

1. normal
2. lognormal
3. uniform
4. loguniform
5. triangular
6. exponential
7. beta
8. gamma
9. gumbel
10. frechet
11. weibull
12. histogram bin
13. poisson
14. binomial
15. negative binomial
16. geometric
17. hypergeometric
18. histogram point

When additional variable types are activated, they assume uniform distributions, and the ordering is as listed on variables.

Examples

Consider the following random variables, distributions and correlations:

- X_1, normal, uncorrelated with others
- X_2, normal, correlated with X_3, X_4 and X_5
- X_3, weibull , correlated with X_5
- X_4, exponential, correlated with X_3, X_4 and X_5
- X_5, normal, correlated with X_5 These correlations are captured by the following commands (order of the variables is respected).

```
uncertain_correlation_matrix
# ordering normal, exponential, weibull
# \f$X_1\f$ \f$X_2\f$ \f$X_5\f$ \f$X_4$\f \f$X_3\f$
1.00 0.00 0.00 0.00 0.00
0.00 1.00 0.50 0.24 0.78
0.00 0.50 1.00 0.00 0.20
0.00 0.24 0.00 1.00 0.49
0.00 0.78 0.20 0.49 1.0
```

5.4.28 continuous_interval_uncertain

- Keywords Area
- variables
- continuous_interval_uncertain

Epistemic uncertain variable - continuous values within one or more intervals

Topics

This keyword is related to the topics:

- continuous_variables
- epistemic_uncertain_variables

Specification

Alias: interval_uncertain

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_intervals	Specify the number of intervals for each variable

	Optional	interval_-probabilities	Assign probability mass to each interval
	Required	lower_bounds	Specify minimum values
	Required	upper_bounds	Specify maximium values
	Optional	initial_point	Initial values
	Optional	descriptors	Label the variables

Description

In addition to continuous and discrete aleatory probability distributions, Dakota provides support for continuous and discrete epistemic uncertainties through the keywords:

- continuous_interval_uncertain
- discrete_interval_uncertain
- discrete_uncertain_set_integer
- discrete_uncertain_set_real

Interval-based and set variables do not represent probability distributions.

In the case of continuous intervals, they can specify a single interval per variable which may be used in interval analysis, where the goal is to determine the interval bounds on the output corresponding to the interval bounds on the input. All values between the bounds are permissible. More detailed continuous interval representations can specify a set of belief structures based on intervals that may be contiguous, overlapping, or disjoint. This is used in specifying the inputs necessary for an epistemic uncertainty analysis using Dempster-Shafer theory of evidence.

Examples

The following specification is for an interval analysis:

```
continuous_interval_uncertain = 2
 lower_bounds = 2.0 4.0
 upper_bounds = 2.5 5.0
```

The following specification is for a Dempster-Shafer analysis:

```
continuous_interval_uncertain = 2
 num_intervals = 3 2
 interval_probs = 0.25 0.5 0.25 0.4 0.6
 lower_bounds = 2.0 4.0 4.5 1.0 3.0
 upper_bounds = 2.5 5.0 6.0 5.0 5.0
```

Here there are 2 interval uncertain variables. The first one is defined by three intervals, and the second by two intervals. The three intervals for the first variable have basic probability assignments of 0.2, 0.5, and 0.3, respectively, while the basic probability assignments for the two intervals for the second variable are 0.4 and 0.6. The basic probability assignments for each interval variable must sum to one. The interval bounds for the first variable are [2, 2.5], [4, 5], and [4.5, 6], and the interval bounds for the second variable are [1.0, 5.0] and [3.0, 5.0]. Note that the intervals can be overlapping or disjoint. The BPA for the first variable indicates that it is twice as likely that the value occurs on the interval [4,5] than either [2,2.5] or [4.5,6].

Theory

The continuous interval uncertain variable is NOT a probability distribution. Although it may seem similar to a histogram, the interpretation of this uncertain variable is different. It is used in epistemic uncertainty analysis, where one is trying to model uncertainty due to lack of knowledge. The continuous interval uncertain variable is used in both interval analysis and in Dempster-Shafer theory of evidence.

- interval analysis -only one interval is allowed for each `continuous_interval_uncertain` variable -the interval is defined by lower and upper bounds -the value of the random variable lies somewhere in this interval -output is the minimum and maximum function value conditional on the specified interval

- Dempster-Shafer theory of evidence -multiple intervals can be assigned to each `continuous_interval_uncertain` variable -a Basic Probability Assignment (BPA) is associated with each interval. The BPA represents a probability that the value of the uncertain variable is located within that interval. -each interval is defined by lower and upper bounds -outputs are called "belief" and "plausibility." Belief represents the smallest possible probability that is consistent with the evidence, while plausibility represents the largest possible probability that is consistent with the evidence. Evidence is the intervals together with their BPA.

num_intervals

- Keywords Area
- variables
- continuous_interval_uncertain
- num_intervals

Specify the number of intervals for each variable

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: iuv_num_intervals
Argument(s): INTEGERLIST

Description

In Dakota, epistemic uncertainty analysis is performed using either interval estimation or Dempster-Shafer theory of evidence. In these approaches, one does not assign a probability distribution to each uncertain input variable. Rather, one divides each uncertain input variable into one or more intervals. The input parameters are only known to occur within intervals; nothing more is assumed.

interval_probabilities

- Keywords Area
- variables
- continuous_interval_uncertain
- interval_probabilities

Assign probability mass to each interval

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: interval_probs iuv_interval_probs
Argument(s): REALLIST

Description

The basic probability assignments for each interval variable must sum to one.

lower_bounds

- Keywords Area
- variables
- continuous_interval_uncertain
- lower_bounds

Specify minimum values

Specification

Alias: none
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- continuous_interval_uncertain
- upper_bounds

Specify maximium values

Specification

Alias: none
Argument(s): REALLIST

Description

Specify maximium values

initial_point

- Keywords Area
- variables
- continuous_interval_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- continuous_interval_uncertain
- descriptors

Label the variables

Specification

Alias: iuv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.29 discrete_interval_uncertain

- Keywords Area
- variables
- discrete_interval_uncertain

Epistemic uncertain variable - one or more intervals

Topics

This keyword is related to the topics:

- continuous_variables
- epistemic_uncertain_variables

Specification

Alias: discrete_uncertain_range
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_intervals	Specify the number of intervals for each variable
	Optional		interval_-probabilities	Assign probability mass to each interval
	Required		lower_bounds	Specify minimum values
	Required		upper_bounds	Specify maximium values
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

In addition to continuous and discrete aleatory probability distributions, Dakota provides support for continuous and discrete epistemic uncertainties through the keywords:

- continuous_interval_uncertain
- discrete_interval_uncertain
- discrete_uncertain_set_integer
- discrete_uncertain_set_real

Interval-based and set variables do not represent probability distributions.

In the case of discrete intervals, they can specify a single interval per variable which may be used in interval analysis, where the goal is to determine the interval bounds on the output corresponding to the interval bounds on the input. Permissible values are any integer within the bound. More detailed continuous interval representations can specify a set of belief structures based on intervals that may be contiguous, overlapping, or disjoint. This is used in specifying the inputs necessary for an epistemic uncertainty analysis using Dempster-Shafer theory of evidence.

Note that the intervals can be overlapping or disjoint.

Examples

Let d1 be 2, 3 or 4 with probability 0.2, 4 or 5 with probability 0.5 and 6 with probability 0.3. Let d2 be 4, 5 or 6 with probability 0.4 and 6, 7 or 8 with probability 0.6. The following specification is for a Dempster-Shafer analysis:

```
discrete_interval_uncertain = 2
 num_intervals = 3 2
 interval_probs = 0.2 0.5 0.3 0.4 0.6
 lower_bounds = 2 4 6 4 6
 upper_bounds = 4 5 6 6 8
```

Theory

- Dempster-Shafer theory of evidence -multiple intervals can be assigned to each `discrete_interval_uncertain` variable -a Basic Probability Assignment (BPA) is associated with each interval. The BPA represents a probability that the value of the uncertain variable is located within that interval. -each interval is defined by lower and upper bounds -outputs are called ”belief” and ”plausibility.” Belief represents the smallest possible probability that is consistent with the evidence, while plausibility represents the largest possible probability that is consistent with the evidence. Evidence is the intervals together with their BPA.

num_intervals

- Keywords Area
- variables
- discrete_interval_uncertain
- num_intervals

Specify the number of intervals for each variable

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGERLIST

Description

In Dakota, epistemic uncertainty analysis is performed using either interval estimation or Dempster-Shafer theory of evidence. In these approaches, one does not assign a probability distribution to each uncertain input variable. Rather, one divides each uncertain input variable into one or more intervals. The input parameters are only known to occur within intervals; nothing more is assumed.

interval_probabilities

- Keywords Area
- variables
- discrete_interval_uncertain
- interval_probabilities

Assign probability mass to each interval

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: interval_probs range_probabilities range_probs
Argument(s): REALLIST

Description

The basic probability assignments for each interval variable must sum to one.

lower_bounds

- Keywords Area
- variables
- discrete_interval_uncertain
- lower_bounds

Specify minimum values

Specification

Alias: none
Argument(s): INTEGERLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- discrete_interval_uncertain
- upper_bounds

Specify maximium values

Specification

Alias: none
Argument(s): INTEGERLIST

Description

Specify maximium values

initial_point

- Keywords Area
- variables
- discrete_interval_uncertain
- initial_point

Initial values

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- discrete_interval_uncertain
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.30 discrete_uncertain_set_integer

- Keywords Area
- variables
- discrete_uncertain_set_integer

Discrete, epistemic uncertain variable - integers within a set

Topics

This keyword is related to the topics:

- discrete_variables
- epistemic_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_set_values	The size of each variables' set of permissible values

	Required	set_values	The permissible values for each discrete variable
	Optional	set_probabilities	
	Optional	initial_point	Initial values
	Optional	descriptors	Label the variables

Description

In addition to continuous and discrete aleatory probability distributions, Dakota provides support for continuous and discrete epistemic uncertainties through the keywords:

- continuous_interval_uncertain
- discrete_interval_uncertain
- discrete_uncertain_set_integer
- discrete_uncertain_set_real

Interval-based and set variables do not represent probability distributions.

Discrete set variables may be used to specify categorical choices which are epistemic. For example, if we have three possible forms for a physics model (model 1, 2, or 3) and there is epistemic uncertainty about which one is correct, a discrete uncertain set may be used to represent this type of uncertainty.

This variable is defined by a set of integers, in which the discrete value may take any value within the integer set (for example, the set may be defined as 1, 2, and 4)

Examples

Let d1 be 2 or 13 and d2 be 4, 5 or 26. The following specification is for an interval analysis:

```
discrete_uncertain_set_integer
 num_set_values 2     3
 set_values     2 13  4 5 26
 descriptors    'd1'  'd2'
```

Theory

The `discrete_uncertain_set_integer` variable is NOT a discrete random variable. It can be contrasted to a the histogram-defined random variables: histogram_bin_uncertain and histogram_point_uncertain. It is used in epistemic uncertainty analysis, where one is trying to model uncertainty due to lack of knowledge.

The discrete uncertain set integer variable is used in both interval analysis and in Dempster-Shafer theory of evidence.

- interval analysis -the values are integers, equally weighted -the true value of the random variable is one of the integers in this set -output is the minimum and maximum function value conditional on the specified inputs
- Dempster-Shafer theory of evidence -the values are integers, but they can be assigned different weights -outputs are called "belief" and "plausibility." Belief represents the smallest possible probability that is consistent with the evidence, while plausibility represents the largest possible probability that is consistent with the evidence. Evidence is the values together with their weights.

num_set_values

- Keywords Area
- variables
- discrete_uncertain_set_integer
- num_set_values

The size of each variables' set of permissible values

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGERLIST

Description

For each variable, the number of permissible values must be specified.

set_values

- Keywords Area
- variables
- discrete_uncertain_set_integer
- set_values

The permissible values for each discrete variable

Specification

Alias: none
Argument(s): INTEGERLIST

Description

This keyword allows specification of the permissible values for discrete set variables. All the variables of the type use this single keyword, so the specification can be confusing. See the description on the discrete_variables page.

set_probabilities

- Keywords Area
- variables
- discrete_uncertain_set_integer
- set_probabilities

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: set_probs
Argument(s): REALLIST

initial_point

- Keywords Area
- variables
- discrete_uncertain_set_integer
- initial_point

Initial values

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- discrete_uncertain_set_integer
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.31 discrete_uncertain_set_real

- Keywords Area
- variables
- discrete_uncertain_set_real

Discrete, epistemic uncertain variable - real numbers within a set

Topics

This keyword is related to the topics:

- discrete_variables
- epistemic_uncertain_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_set_values	The size of each variables' set of permissible values
	Required		set_values	The permissible values for each discrete variable
	Optional		set_probabilities	
	Optional		initial_point	Initial values
	Optional		descriptors	Label the variables

Description

In addition to continuous and discrete aleatory probability distributions, Dakota provides support for continuous and discrete epistemic uncertainties through the keywords:

- continuous_interval_uncertain
- discrete_interval_uncertain

- discrete_uncertain_set_integer
- discrete_uncertain_set_real

Interval-based and set variables do not represent probability distributions.

Discrete set variables may be used to specify categorical choices which are epistemic. For example, if we have three possible forms for a physics model (model 1, 2, or 3) and there is epistemic uncertainty about which one is correct, a discrete uncertain set may be used to represent this type of uncertainty.

This variable is defined by a set of reals, in which the discrete variable may take any value defined within the real set (for example, a parameter may have two allowable real values, 3.285 or 4.79).

Examples

Let d1 be 2.1 or 1.3 and d2 be 0.4, 5 or 2.6. The following specification is for an interval analysis:

```
discrete_uncertain_set_integer
 num_set_values  2           3
 set_values      2.1  1.3    0.4  5  2.6
 descriptors     'd1'        'd2'
```

Theory

The `discrete_uncertain_set_integer` variable is NOT a discrete random variable. It can be contrasted to a the histogram-defined random variables: histogram_bin_uncertain and histogram_point_uncertain. It is used in epistemic uncertainty analysis, where one is trying to model uncertainty due to lack of knowledge.

The discrete uncertain set integer variable is used in both interval analysis and in Dempster-Shafer theory of evidence.

- interval analysis -the values are integers, equally weighted -the true value of the random variable is one of the integers in this set -output is the minimum and maximum function value conditional on the specified inputs
- Dempster-Shafer theory of evidence -the values are integers, but they can be assigned different weights -outputs are called ”belief” and ”plausibility.” Belief represents the smallest possible probability that is consistent with the evidence, while plausibility represents the largest possible probability that is consistent with the evidence. Evidence is the values together with their weights.

num_set_values

- Keywords Area
- variables
- discrete_uncertain_set_real
- num_set_values

The size of each variables’ set of permissible values

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGERLIST

Description

For each variable, the number of permissible values must be specified.

set_values

- Keywords Area
- variables
- discrete_uncertain_set_real
- set_values

The permissible values for each discrete variable

Specification

Alias: none
Argument(s): REALLIST

Description

This keyword allows specification of the permissible values for discrete set variables. All the variables of the type use this single keyword, so the specification can be confusing. See the description on the discrete_variables page.

set_probabilities

- Keywords Area
- variables
- discrete_uncertain_set_real
- set_probabilities

TODO: Need a Blurb

Topics

This keyword is related to the topics:

- empty

Specification

Alias: set_probs
Argument(s): REALLIST

initial_point

- Keywords Area
- variables
- discrete_uncertain_set_real
- initial_point

Initial values

Specification

Alias: none
Argument(s): REALLIST

Description

The `initial_point` specifications provide the point in design space from which an iterator is started and defaults to the midpoint of bounds (continuous design variables) or the middle value (discrete design variables).

descriptors

- Keywords Area
- variables
- discrete_uncertain_set_real
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.32 continuous_state

- Keywords Area
- variables
- continuous_state

Continuous state variables

Topics

This keyword is related to the topics:

- state_variables
- continuous_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_state	Initial values for the state variables
	Optional		lower_bounds	Specify minimum values
	Optional		upper_bounds	Specify maximium values
	Optional		descriptors	Label the variables

Description

Continuous state variables are defined by bounds.

Default behavior for most methods is that only the initial_state values are used.

See the state_variables page for details on the behavior of state variables.

initial_state

- Keywords Area
- variables
- continuous_state
- initial_state

Initial values for the state variables

Specification

Alias: csv_initial_state

Argument(s): REALLIST

Description

The `initial_state` specifications provide the initial values for the state variables.

This is an optional keyword. If it is not specified, the initial state will be inferred from the other keywords that define the state variable.

Defaults are:

- Continuous state variables - use the midpoint of the bounds

- Set variables - use the value with the index closest to the middle of the set
- Range variables - use the midpoint of the range

lower_bounds

- Keywords Area
- variables
- continuous_state
- lower_bounds

Specify minimum values

Specification

Alias: csv_lower_bounds
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- continuous_state
- upper_bounds

Specify maximium values

Specification

Alias: csv_upper_bounds
Argument(s): REALLIST

Description

Specify maximium values

descriptors

- Keywords Area
- variables
- continuous_state
- descriptors

Label the variables

Specification

Alias: csv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.33 discrete_state_range

- Keywords Area
- variables
- discrete_state_range

Discrete state variable defined by bounds

Topics

This keyword is related to the topics:

- discrete_variables
- state_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_state	Initial values for the state variables
	Optional		lower_bounds	Specify minimum values
	Optional		upper_bounds	Specify maximium values
	Optional		descriptors	Label the variables

Description

Discrete state variables defined by bounds.

The details of how to specify this discrete variable are located on the discrete_variables page.

See the state_variables page for details on the behavior of state variables.

initial_state

- Keywords Area
- variables
- discrete_state_range
- initial_state

Initial values for the state variables

Specification

Alias: dsv_initial_state
Argument(s): INTEGERLIST

Description

The `initial_state` specifications provide the initial values for the state variables.

This is an optional keyword. If it is not specified, the initial state will be inferred from the other keywords that define the state variable.

Defaults are:

- Continuous state variables - use the midpoint of the bounds
- Set variables - use the value with the index closest to the middle of the set
- Range variables - use the midpoint of the range

lower_bounds

- Keywords Area
- variables
- discrete_state_range
- lower_bounds

Specify minimum values

Specification

Alias: dsv_lower_bounds
Argument(s): INTEGERLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- variables
- discrete_state_range
- upper_bounds

Specify maximium values

Specification

Alias: dsv_upper_bounds
Argument(s): INTEGERLIST

Description

Specify maximium values

descriptors

- Keywords Area
- variables
- discrete_state_range
- descriptors

Label the variables

Specification

Alias: dsv_descriptors
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.34 discrete_state_set_integer

- Keywords Area
- variables
- discrete_state_set_integer

Discrete state variables, each defined by a set of permissible integers

Topics

This keyword is related to the topics:

- discrete_variables
- state_variables

Specification

Alias: none

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_state	Initial values for the state variables
	Optional		num_set_values	The size of each variables' set of permissible values
	Required		set_values	The permissible values for each discrete variable
	Optional		descriptors	Label the variables

Description

Discrete state variables defined by a set of permissible integers.

The details of how to specify this discrete variable are located on the discrete_variables page.

See the state_variables page for details on the behavior of state variables.

initial_state

- Keywords Area
- variables
- discrete_state_set_integer
- initial_state

Initial values for the state variables

Specification

Alias: none

Argument(s): INTEGERLIST

Description

The `initial_state` specifications provide the initial values for the state variables.

This is an optional keyword. If it is not specified, the initial state will be inferred from the other keywords that define the state variable.

Defaults are:

- Continuous state variables - use the midpoint of the bounds
- Set variables - use the value with the index closest to the middle of the set
- Range variables - use the midpoint of the range

num_set_values

- Keywords Area
- variables
- discrete_state_set_integer
- num_set_values

The size of each variables' set of permissible values

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none

Argument(s): INTEGERLIST

Description

For each variable, the number of permissible values must be specified.

set_values

- Keywords Area
- variables
- discrete_state_set_integer
- set_values

The permissible values for each discrete variable

Specification

Alias: none

Argument(s): INTEGERLIST

Description

This keyword allows specification of the permissible values for discrete set variables. All the variables of the type use this single keyword, so the specification can be confusing. See the description on the discrete_variables page.

descriptors

- Keywords Area
- variables
- discrete_state_set_integer
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.4.35 discrete_state_set_real

- Keywords Area
- variables
- discrete_state_set_real

Discrete state variables, each defined by a set of permissible real numbers

Topics

This keyword is related to the topics:

- discrete_variables
- state_variables

Specification

Alias: none
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		initial_state	Initial values for the state variables
	Optional		num_set_values	The size of each variables' set of permissible values
	Required		set_values	The permissible values for each discrete variable
	Optional		descriptors	Label the variables

Description

Discrete state variables defined by a set of permissible real numbers.

The details of how to specify this discrete variable are located on the discrete_variables page.

See the state_variables page for details on the behavior of state variables.

initial_state

- Keywords Area
- variables
- discrete_state_set_real
- initial_state

Initial values for the state variables

Specification

Alias: none

Argument(s): REALLIST

Description

The `initial_state` specifications provide the initial values for the state variables.

This is an optional keyword. If it is not specified, the initial state will be inferred from the other keywords that define the state variable.

Defaults are:

- Continuous state variables - use the midpoint of the bounds
- Set variables - use the value with the index closest to the middle of the set
- Range variables - use the midpoint of the range

num_set_values

- Keywords Area
- variables
- discrete_state_set_real
- num_set_values

The size of each variables' set of permissible values

Topics

This keyword is related to the topics:

- not_yet_reviewed

Specification

Alias: none
Argument(s): INTEGERLIST

Description

For each variable, the number of permissible values must be specified.

set_values

- Keywords Area
- variables
- discrete_state_set_real
- set_values

The permissible values for each discrete variable

Specification

Alias: none
Argument(s): REALLIST

Description

This keyword allows specification of the permissible values for discrete set variables. All the variables of the type use this single keyword, so the specification can be confusing. See the description on the discrete_variables page.

descriptors

- Keywords Area
- variables
- discrete_state_set_real
- descriptors

Label the variables

Specification

Alias: none
Argument(s): STRINGLIST

Description

The optional variables labels specification `descriptors` is a list of strings which will be printed in Dakota input file to identify the values for particular variable.

The default descriptor strings use a root string plus a numeric identifier.

5.5 interface

- Keywords Area
- interface

Specifies how function evaluations will be performed in order to map the variables into the responses.

Topics

This keyword is related to the topics:

- block
- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		id_interface	Give the block a name, in case of multiple interface blocks

	Optional	algebraic_-mappings	Use AMPL to define algebraic input-output mappings
	Optional	analysis_drivers	Define how Dakota should run a function evaluation
	Optional	asynchronous	Specify analysis driver concurrency, when Dakota is run in serial
	Optional	evaluation_servers	Specify the number of evaluation servers when Dakota is run in parallel
	Optional	evaluation_-scheduling	Specify the scheduling of concurrent evaluations when Dakota is run in parallel
	Optional	processors_per_-evaluation	Specify the number of processors per evaluation server when Dakota is run in parallel
	Optional	analysis_servers	Specify the number of analysis servers when Dakota is run in parallel
	Optional	analysis_-scheduling	Specify the scheduling of concurrent analyses when Dakota is run in parallel

Description

The interface section in a Dakota input file specifies how function evaluations will be performed in order to map the variables into the responses.

In this context, a ”function evaluation” is the series of operations that takes the variables and computes the responses. This can be comprised of one or many codes, scripts, and glue, which are generically referred to as ”analysis drivers”.

The optional `asynchronous` flag specifies use of asynchronous protocols (i.e., background system calls, nonblocking forks, POSIX threads) when evaluations or analyses are invoked. The `evaluation_concurrency` and `analysis_concurrency` specifications serve a dual purpose:

- when running Dakota on a single processor in `asynchronous` mode, the default concurrency of evaluations and analyses is all concurrency that is available. The `evaluation_concurrency` and `analysis_concurrency` specifications can be used to limit this concurrency in order to avoid machine overload or usage policy violation.

- when running Dakota on multiple processors in message passing mode, the default concurrency of evaluations and analyses on each of the servers is one (i.e., the parallelism is exclusively that of the message passing). With the `evaluation_concurrency` and `analysis_concurrency` specifications, a hybrid parallelism can be selected through combination of message passing parallelism with asynchronous parallelism on each server.

If Dakota's automatic parallel configuration is undesirable for some reason, the user can specify overrides that enforce a desired number of partitions, a size for the partitions, and/or a desired scheduling configuration at the evaluation and analysis parallelism levels. The optional `evaluation_servers` and `analysis_servers` specifications support user overrides of the automatic parallel configuration for the number of evaluation servers and the number of analysis servers, and the optional `processors_per_evaluation` specification supports user overrides for the size of processor allocations for evaluation servers (Note: see direct for the `processors_per_analysis` specification supported for direct interfaces). Similarly, the optional `evaluation_scheduling` and `analysis_scheduling` specifications can be used to override the automatic parallel configuration at the evaluation and analysis parallelism levels to use either a dedicated `master` or a `peer` partition. In addition, the evaluation parallelism level supports an override for the scheduling algorithm used within a `peer` partition; this can be either `dynamic` or `static` scheduling (default configuration of a peer partition employs a dynamic scheduler when it can be supported; i.e., when the peer 1 local scheduling can be asynchronous). The ParallelLibrary class and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] provide additional details on parallel configurations.

When performing asynchronous local evaluations, the `local_evaluation_scheduling` keyword controls how new evaluation jobs are dispatched when one completes. If the `local_evaluation_scheduling` is specified as `dynamic` (the default), each completed evaluation will be replaced by the next in the local evaluation queue. If `local_evaluation_scheduling` is specified as `static`, each completed evaluation will be replaced by an evaluation number that is congruent modulo the `evaluation_concurrency`. This is helpful for relative node scheduling as described in `Dakota/examples/parallelism`. For example, assuming only asynchronous local concurrency (no MPI), if the local concurrency is 6 and job 2 completes, it will be replaced with job 8. For the case of hybrid parallelism, static local scheduling results in evaluation replacements that are modulo the total capacity, defined as the product of the evaluation concurrency and the number of evaluation servers. Both of these cases can result in idle processors if runtimes are non-uniform, so the default dynamic scheduling is preferred when relative node scheduling is not required.

Theory

Function evaluations are performed using either interfaces to simulation codes, algebraic mappings, or a combination of the two.

When employing mappings with simulation codes, the interface invokes the simulation using either forks, direct function invocations, or computational grid invocations.

- In the fork case, Dakota will treat the simulation as a black-box and communication between Dakota and the simulation occurs through parameter and result files. This is the most common case.

- In the direct function case, the simulation is internal to Dakota and communication occurs through the function parameter list. The direct case can involve linked simulation codes or test functions which are compiled into the Dakota executable. The test functions allow for rapid testing of algorithms without process creation overhead or engineering simulation expense.

- The grid case is experimental and under development, but is intended to support simulations which are external to Dakota and geographically distributed.

When employing algebraic mappings, the AMPL solver library [[28] "Gay, 1997"] is used to evaluate a directed acyclic graph (DAG) specification from a separate stub.nl file. Separate stub.col and stub.row files are also required to declare the string identifiers of the subset of inputs and outputs, respectively, that will be used in the algebraic mappings.

5.5.1 id_interface

- Keywords Area
- interface
- id_interface

Give the block a name, in case of multiple interface blocks

Topics

This keyword is related to the topics:

- block_identifier

Specification

Alias: none
Argument(s): STRING

Description

The optional set identifier specification uses the keyword `id_interface` to input a string for use in identifying a particular interface specification. A model can then identify the use of this interface by specifying the same string in its `interface_pointer` specification.

If the `id_interface` specification is omitted, a particular interface specification will be used by a model only if that model omits specifying a `interface_pointer` and if the interface set was the last set parsed (or is the only set parsed). In common practice, if only one interface set exists, then `id_interface` can be safely omitted from the interface specification and `interface_pointer` can be omitted from the model specification(s), since there is no potential for ambiguity in this case.

Examples

For example, a model whose specification contains `interface_pointer` = 'I1' will use an interface specification with `id_interface` = 'I1'.

5.5.2 algebraic_mappings

- Keywords Area
- interface
- algebraic_mappings

Use AMPL to define algebraic input-output mappings

Specification

Alias: none
Argument(s): STRING

Description

If desired, one can define algebraic input-output mappings using the AMPL code [[25] "Fourer et al., 2003"] and save these mappings in 3 files: `stub.nl`, `stub.col`, and `stub.row`, where `stub` is a particular root name describing a particular problem. These files names can be communicated to Dakota using the `algebraic_mappings` input. This string may either specify the `stub.nl` filename, or alternatively, just the `stub` itself.

Dakota then uses `stub.col` and `stub.row` to extract the input and output identifier strings and employs the AMPL solver library [[28] "Gay, 1997"] to process the DAG specification in `stub.nl`. The variable and objective function names declared within AMPL should be a subset of the variable descriptors and response descriptors used by Dakota (see variables and descriptors). Ordering is not important, as Dakota will reorder data as needed.

5.5.3 analysis_drivers

- Keywords Area
- interface
- analysis_drivers

Define how Dakota should run a function evaluation

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): STRINGLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		analysis_-components	Provide additional identifiers to analysis drivers.
	Optional		input_filter	Run a pre-processing script before the analysis drivers

	Optional		output_filter	Run a post-processing script after the analysis drivers (Not recommended)
	Required*(Choose One)*	**Group 1**	system	Launch analysis drivers with a system call
			fork	Launch analysis drivers using fork command
			direct	Run analysis drivers that are linked-to or compiled-with Dakota
			matlab	Run Matlab with a direct interface - requires special Dakota build
			python	Run Python with a direct interface - requires special Dakota build
			scilab	Run Scilab with a direct interface - requires special Dakota build
			grid	Experimental capability
	Optional		failure_capture	Determine how Dakota responds to analysis driver failure
	Optional		deactivate	Deactivate Dakota features to simplify interface development, increase execution speed, or reduce memory and disk requirements

Description

The required `analysis_drivers` keyword provides the names of executable analysis programs or scripts, a.k.a. ”drivers” which comprise a function evaluation. The optional and required sub-keywords specify how Dakota will manage directories and files, and run the driver(s).

Types of Interfaces

Dakota has two recommended ways of running analysis drivers:

- as an external processes (`fork`), or
- using internal code to couple to the analysis driver (`direct`)

Other options are available for advanced users, and are not as well documented, supported, or tested:

- external processes (`system`)
- internal coupling (`python`, `matlab`, `scilab`, `grid`)

Use Cases

The internally coupled codes have few options because many of the details are already handled with the coupling. Their behavior is described in the direct keyword.

For external processes using the fork keyword,

A function evaluation may comprise:

1. *A single analysis driver*: Function evaluation, including all pre- and post-processing is contained entirely within a single script/executable.
2. *A single analysis driver with filters*: Function evaluation is explicitly split into pre-processing (performed by the input filter), analysis, and post-processing (by the output filter).
3. *A single analysis driver with environment variables*: Function evaluation is contained within one analysis driver, but it requires environment variables to be set before running.
4. *Multiple analysis drivers*: Drivers are run sequentially or concurrently (See the `asynchronous` keyword) and can have any of the above options as well.

For details and examples see the Simulation Interface Components section of the Interfaces chapter of the User's Manual; for details on the filters and environment variables, see the subsection on Syntax for Filter and Driver Strings.

analysis_components

- Keywords Area
- interface
- analysis_drivers
- analysis_components

Provide additional identifiers to analysis drivers.

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none

Argument(s): STRINGLIST

Description

The optional `analysis_components` specification allows the user to provide additional identifiers (e.g., mesh file names) for use by the analysis drivers. This is particularly useful when the same analysis driver is to be reused multiple times for slightly different analyses. The specific content within the strings is open-ended and can involve whatever syntax is convenient for a particular analysis driver. The number of analysis components n_c should be an integer multiple of the number of drivers n_d, and the first n_c/n_d component strings will be passed to the first driver, etc.

input_filter

- Keywords Area
- interface
- analysis_drivers
- input_filter

Run a pre-processing script before the analysis drivers

Specification

Alias: none
Argument(s): STRING

Description

The optional `input_filter` and `output_filter` specifications provide the names of separate pre- and post-processing programs or scripts which assist in mapping Dakota parameters files into analysis input files and mapping analysis output files into Dakota results files, respectively.

If there is only a single analysis driver, then it is usually most convenient to combine pre- and post-processing requirements into a single analysis driver script and omit the separate input and output filters. However, in the case of multiple analysis drivers, the input and output filters provide a convenient location for non-repeated pre- and post-processing requirements. That is, input and output filters are only executed once per function evaluation, regardless of the number of analysis drivers, which makes them convenient locations for data processing operations that are shared among the analysis drivers.

output_filter

- Keywords Area
- interface
- analysis_drivers
- output_filter

Run a post-processing script after the analysis drivers

Specification

Alias: none
Argument(s): STRING

Description

The optional `input_filter` and `output_filter` specifications provide the names of separate pre- and post-processing programs or scripts which assist in mapping Dakota parameters files into analysis input files and mapping analysis output files into Dakota results files, respectively.

If there is only a single analysis driver, then it is usually most convenient to combine pre- and post-processing requirements into a single analysis driver script and omit the separate input and output filters. However, in the case of multiple analysis drivers, the input and output filters provide a convenient location for non-repeated pre- and post-processing requirements. That is, input and output filters are only executed once per function evaluation, regardless of the number of analysis drivers, which makes them convenient locations for data processing operations that are shared among the analysis drivers.

system

- Keywords Area
- interface
- analysis_drivers
- system

(Not recommended) Launch analysis drivers with a system call

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		parameters_file	Specify the name of the parameters file
	Optional		results_file	Specify the name of the results file
	Optional		allow_existing_-results	Change how Dakota deals with existing results files
	Optional		verbatim	Specify the command Dakota uses to launch analysis driver(s) and filters

	Optional	aprepro	Write parameters files in APREPRO syntax
	Optional	file_tag	Tag each parameters & results file name with the function evaluation number
	Optional	file_save	Keep the parameters & results files after the analysis driver completes
	Optional	work_directory	Have Dakota set up working directories before launching analysis driver

Description

The system call interface is included in Dakota for portability and backward compatibility. Users are strongly encouraged to use the `fork` interface if possible, reverting to system only when necessary. To enable the system call interface, replace the `fork` keyword with `system`. All other keywords have identical meanings to those for the fork interface

See Also

These keywords may also be of interest:

- fork

parameters_file

- Keywords Area
- interface
- analysis_drivers
- system
- parameters_file

Specify the name of the parameters file

Specification

Alias: none
Argument(s): STRING

Description

The parameters file is used by Dakota to pass the parameter values to the analysis driver. The name of the file can be optionally specified using the `parameters_file` keyword.

If this is not specified, the default data transfer files are temporary files with system-generated names (e.g., `/usr/tmp/aaaa08861`).

results_file

- Keywords Area
- interface
- analysis_drivers
- system
- results_file

Specify the name of the results file

Specification

Alias: none

Argument(s): STRING

Description

The results file must be written by the analysis driver. It is read by Dakota to determine the response values for each function evaluation.

The name of the file can be optionally specified using the `results_file` keyword.

If this is not specified, the default data transfer files are temporary files with system-generated names (e.g., `/usr/tmp/aaaa08861`).

allow_existing_results

- Keywords Area
- interface
- analysis_drivers
- system
- allow_existing_results

Change how Dakota deals with existing results files

Specification

Alias: none

Argument(s): none

Description

By default Dakota will remove existing results files before invoking the `analysis_driver` to avoid problems created by stale files in the current directory. To override this behavior and not delete existing files, specify `allow_existing_results`.

verbatim

- Keywords Area
- interface
- analysis_drivers
- system
- verbatim

Specify the command Dakota uses to launch analysis driver(s) and filters

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none

Argument(s): none

Description

The typical commands that Dakota uses to launch analysis drivers are:

```
> input_filter parameters_file_name results_file_name
> analysis_driver parameters_file_name results_file_name
> output_filter parameters_file_name results_file_name
```

Dakota will automatically arrange the executables and file names.

If the analysis driver requires a different syntax, the entire command can be specified as the analysis driver and the `verbatim` keyword will tell Dakota to use this as the command.

Note, this will not allow the use of `file_tag`, because the full command must be specified.

For additional information on invocation syntax, see the Interfaces chapter of the Users Manual [[4] "Adams et al., 2010"].

aprepro

- Keywords Area
- interface
- analysis_drivers
- system
- aprepro

Write parameters files in APREPRO syntax

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: dprepro
Argument(s): none

Description

The format of data in the parameters files can be modified for direct usage with the APREPRO pre-processing tool [[76] "Sjaardema, 1992"] using the `aprepro` specification

Without this keyword, the parameters file are written in DPrePro format. DPrePro is a utility included with Dakota, described in the Users Manual [[4] "Adams et al., 2010"].

file_tag

- Keywords Area
- interface
- analysis_drivers
- system
- file_tag

Tag each parameters & results file name with the function evaluation number

Specification

Alias: none
Argument(s): none

Description

If this keyword is used, Dakota will append a period and the function evaluation number to the names of the parameter and results files.

For example, if the following is included in the `interface` section of the Dakota input:

```
parameters_file = params.in
results_file = results.out
file_tag
```

Then for the 3rd evaluation, Dakota will write `params.in.3`, and will expect `results.out.3` to be written by the analysis driver.

If this keyword is omitted, the default is no file tagging.

File tagging is most useful when multiple function evaluations are running simultaneously using files in a shared disk space. The analysis driver will be able to infer the function evaluation number from the file names. Note that when the `file_save` keyword is used, Dakota renames parameters and results files, giving them tags, after execution of the analysis driver if they otherwise would be overwritten by the next evaluation.

file_save

- Keywords Area
- interface
- analysis_drivers
- system
- file_save

Keep the parameters & results files after the analysis driver completes

Specification

Alias: none
Argument(s): none

Description

If `file_save` is used, Dakota will not delete the parameters and results files after the function evaluation is completed.

The default behavior is NOT to save these files.

If `file_tag` is not specified and the saved files would be overwritten by a future evaluation, Dakota renames them after the analysis driver has run by tagging them with the evaluation number.

File saving is most useful when debugging the data communication between Dakota and the simulation.

work_directory

- Keywords Area
- interface
- analysis_drivers
- system
- work_directory

Have Dakota set up working directories before launching analysis driver

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		named	Specify the name of the work directory created by Dakota
	Optional		directory_tag	Tag each work directory with the function evaluation number
	Optional		directory_save	Do not delete the work directories after completion of the function evaluation
	Optional*(Choose One)*	**Group 1**	template_directory	Populate each work directory created by Dakota with the contents of a template directory
			template_files	Populate each work directory created by Dakota with the specified files

Description

When performing concurrent evaluations and/or analyses, it is often necessary to cloister input and output files in separate directories to avoid conflicts. When the `work_directory` feature is enabled, Dakota will create a directory for each evaluation/analysis (with optional tagging and saving as with files), and execute the analysis driver from that working directory.

The directory may be `named` with a string, or left anonymous to use temporary file space. If an optional `template_directory` and/or list of `template_files` is specified, Dakota will link (or copy) those files into each working directory.

named

- Keywords Area
- interface
- analysis_drivers
- system
- work_directory
- named

Specify the name of the work directory created by Dakota

Specification

Alias: none
Argument(s): STRING

Description

The `named` keyword allows the user to specify the name of the work directory created by Dakota. The directory will be created as a subdirectory of the directory in which Dakota is invoked.

If `named` is not used, the default work directory is a temporary file with a system-generated name (e.g., `/tmp/file93vb1z/`).

See Also

These keywords may also be of interest:

- directory_tag
- directory_save

directory_tag

- Keywords Area
- interface
- analysis_drivers
- system
- work_directory
- directory_tag

Tag each work directory with the function evaluation number

Specification

Alias: dir_tag
Argument(s): none

Description

If this keyword is used, Dakota will append a period and the function evaluation number to the work directory names.

If this keyword is omitted, the default is no tagging, and the same work directory will be used for ALL function evaluations. Tagging is most useful when multiple function evaluations are running simultaneously.

directory_save

- Keywords Area
- interface
- analysis_drivers
- system
- work_directory
- directory_save

Do not delete the work directories after completion of the function evaluation

Specification

Alias: dir_save
Argument(s): none

Description

If a working directory is created by Dakota using the `work_directory` keyword, it is deleted after the evaluation is completed unless the the `directory_save` keyword is specified.

This keyword can only be used if the `named` keyword is used to specify a work directory name.

template_directory

- Keywords Area
- interface
- analysis_drivers
- system
- work_directory
- template_directory

Populate each work directory created by Dakota with the contents of a template directory

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		copy	Copy template files & directories instead of linking

	Optional	replace	Overwrite existing files within a work directory

Description

Upon creation of a working directory, all the files in the directory specified by `template_directory` are linked (file system symbolic links) or copied into the working directory.

copy

- Keywords Area
- interface
- analysis_drivers
- system
- work_directory
- template_directory
- copy

Copy template files & directories instead of linking

Specification

Alias: none
Argument(s): none

Description

By default, files & directories specified by the `template_files` and `template_directory` keywords are linked into each work directory created by Dakota.

If the `copy` keyword is specified, these files & directories are copied rather than linked.

replace

- Keywords Area
- interface
- analysis_drivers
- system
- work_directory
- template_directory
- replace

Overwrite existing files within a work directory

Specification

Alias: none
Argument(s): none

Description

By default, Dakota will not overwrite any existing files in a work directory it creates. The `replace` keyword changes this behavior to force overwriting.

template_files

- Keywords Area
- interface
- analysis_drivers
- system
- work_directory
- template_files

Populate each work directory created by Dakota with the specified files

Specification

Alias: none
Argument(s): STRINGLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		copy	Copy template files & directories instead of linking
	Optional		replace	Overwrite existing files within a work directory

Description

Upon creation of a working directory, all the files named after the `template_files` keyword are linked (file system symbolic links) into the working directory.

copy

- Keywords Area
- interface
- analysis_drivers
- system

- work_directory
- template_files
- copy

Copy template files & directories instead of linking

Specification

Alias: none
Argument(s): none

Description

By default, files & directories specified by the `template_files` and `template_directory` keywords are linked into each work directory created by Dakota.

If the `copy` keyword is specified, these files & directories are copied rather than linked.

replace

- Keywords Area
- interface
- analysis_drivers
- system
- work_directory
- template_files
- replace

Overwrite existing files within a work directory

Specification

Alias: none
Argument(s): none

Description

By default, Dakota will not overwrite any existing files in a work directory it creates. The `replace` keyword changes this behavior to force overwriting.

fork

- Keywords Area
- interface
- analysis_drivers
- fork

Launch analysis drivers using fork command

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		parameters_file	Specify the name of the parameters file
	Optional		results_file	Specify the name of the results file
	Optional		allow_existing_-results	Change how Dakota deals with existing results files
	Optional		verbatim	Specify the command Dakota uses to launch analysis driver(s) and filters
	Optional		aprepro	Write parameters files in APREPRO syntax
	Optional		file_tag	Tag each parameters & results file name with the function evaluation number
	Optional		file_save	Keep the parameters & results files after the analysis driver completes
	Optional		work_directory	Have Dakota set up working directories before launching analysis driver

Description

The `fork` interface is the most common means by which Dakota launches a separate application analysis process.

The `fork` interface is recommended over `system` for most analysis drivers that are external to Dakota (i.e. not using the `direct` interface).

As explained in the Users Manual, the parameters and results file names are passed on the command line to the analysis driver(s). If input/output filters are specified, they will be run before/after the analysis drivers. The `verbatim` keyword is used to modify the default driver/filter commands.

For additional information on invocation syntax, see the Interfaces chapter of the Users Manual [[4] "Adams et al., 2010"].

Examples

```
interface
  analysis_drivers = 'rosenbrock'
    fork
      parameters_file = 'params.in'
      results_file   = 'results.out'
      file_tag
      file_save
```

parameters_file

- Keywords Area
- interface
- analysis_drivers
- fork
- parameters_file

Specify the name of the parameters file

Specification

Alias: none
Argument(s): STRING

Description

The parameters file is used by Dakota to pass the parameter values to the analysis driver. The name of the file can be optionally specified using the `parameters_file` keyword.

If this is not specified, the default data transfer files are temporary files with system-generated names (e.g., `/usr/tmp/aaaa08861`).

results_file

- Keywords Area
- interface
- analysis_drivers
- fork
- results_file

Specify the name of the results file

Specification

Alias: none
Argument(s): STRING

Description

The results file must be written by the analysis driver. It is read by Dakota to determine the response values for each function evaluation.

The name of the file can be optionally specified using the `results_file` keyword.

If this is not specified, the default data transfer files are temporary files with system-generated names (e.g., `/usr/tmp/aaaa08861`).

allow_existing_results

- Keywords Area
- interface
- analysis_drivers
- fork
- allow_existing_results

Change how Dakota deals with existing results files

Specification

Alias: none

Argument(s): none

Description

By default Dakota will remove existing results files before invoking the `analysis_driver` to avoid problems created by stale files in the current directory. To override this behavior and not delete existing files, specify `allow_existing_results`.

verbatim

- Keywords Area
- interface
- analysis_drivers
- fork
- verbatim

Specify the command Dakota uses to launch analysis driver(s) and filters

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): none

Description

The typical commands that Dakota uses to launch analysis drivers are:

```
> input_filter parameters_file_name results_file_name
> analysis_driver parameters_file_name results_file_name
> output_filter parameters_file_name results_file_name
```

Dakota will automatically arrange the executables and file names.

If the analysis driver requires a different syntax, the entire command can be specified as the analysis driver and the `verbatim` keyword will tell Dakota to use this as the command.

Note, this will not allow the use of `file_tag`, because the full command must be specified.

For additional information on invocation syntax, see the Interfaces chapter of the Users Manual [[4] "Adams et al., 2010"].

aprepro

- Keywords Area
- interface
- analysis_drivers
- fork
- aprepro

Write parameters files in APREPRO syntax

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: dprepro
Argument(s): none

Description

The format of data in the parameters files can be modified for direct usage with the APREPRO pre-processing tool [[76] "Sjaardema, 1992"] using the `aprepro` specification

Without this keyword, the parameters file are written in DPrePro format. DPrePro is a utility included with Dakota, described in the Users Manual [[4] "Adams et al., 2010"].

file_tag

- Keywords Area
- interface
- analysis_drivers
- fork
- file_tag

Tag each parameters & results file name with the function evaluation number

Specification

Alias: none
Argument(s): none

Description

If this keyword is used, Dakota will append a period and the function evaluation number to the names of the parameter and results files.

For example, if the following is included in the `interface` section of the Dakota input:

```
parameters_file = params.in
results_file = results.out
file_tag
```

Then for the 3rd evaluation, Dakota will write `params.in.3`, and will expect `results.out.3` to be written by the analysis driver.

If this keyword is omitted, the default is no file tagging.

File tagging is most useful when multiple function evaluations are running simultaneously using files in a shared disk space. The analysis driver will be able to infer the function evaluation number from the file names. Note that when the `file_save` keyword is used, Dakota renames parameters and results files, giving them tags, after execution of the analysis driver if they otherwise would be overwritten by the next evaluation.

file_save

- Keywords Area
- interface
- analysis_drivers
- fork
- file_save

Keep the parameters & results files after the analysis driver completes

Specification

Alias: none
Argument(s): none

Description

If `file_save` is used, Dakota will not delete the parameters and results files after the function evaluation is completed.

The default behavior is NOT to save these files.

If `file_tag` is not specified and the saved files would be overwritten by a future evaluation, Dakota renames them after the analysis driver has run by tagging them with the evaluation number.

File saving is most useful when debugging the data communication between Dakota and the simulation.

work_directory

- Keywords Area
- interface
- analysis_drivers
- fork
- work_directory

Have Dakota set up working directories before launching analysis driver

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		named	Specify the name of the work directory created by Dakota
	Optional		directory_tag	Tag each work directory with the function evaluation number
	Optional		directory_save	Do not delete the work directories after completion of the function evaluation
	Optional(*Choose One*)	**Group 1**	template_directory	Populate each work directory created by Dakota with the contents of a template directory
			template_files	Populate each work directory created by Dakota with the specified files

Description

When performing concurrent evaluations and/or analyses, it is often necessary to cloister input and output files in separate directories to avoid conflicts. When the `work_directory` feature is enabled, Dakota will create a directory for each evaluation/analysis (with optional tagging and saving as with files), and execute the analysis driver from that working directory.

The directory may be `named` with a string, or left anonymous to use temporary file space. If an optional `template_directory` and/or list of `template_files` is specified, Dakota will link (or copy) those files into each working directory.

named

- Keywords Area
- interface
- analysis_drivers
- fork
- work_directory
- named

Specify the name of the work directory created by Dakota

Specification

Alias: none

Argument(s): STRING

Description

The `named` keyword allows the user to specify the name of the work directory created by Dakota. The directory will be created as a subdirectory of the directory in which Dakota is invoked.

If `named` is not used, the default work directory is a temporary file with a system-generated name (e.g., `/tmp/file93vb1z/`).

See Also

These keywords may also be of interest:

- directory_tag
- directory_save

directory_tag

- Keywords Area
- interface
- analysis_drivers

- fork
- work_directory
- directory_tag

Tag each work directory with the function evaluation number

Specification

Alias: dir_tag
Argument(s): none

Description

If this keyword is used, Dakota will append a period and the function evaluation number to the work directory names.

If this keyword is omitted, the default is no tagging, and the same work directory will be used for ALL function evaluations. Tagging is most useful when multiple function evaluations are running simultaneously.

directory_save

- Keywords Area
- interface
- analysis_drivers
- fork
- work_directory
- directory_save

Do not delete the work directories after completion of the function evaluation

Specification

Alias: dir_save
Argument(s): none

Description

If a working directory is created by Dakota using the `work_directory` keyword, it is deleted after the evaluation is completed unless the the `directory_save` keyword is specified.

This keyword can only be used if the `named` keyword is used to specify a work directory name.

template_directory

- Keywords Area
- interface
- analysis_drivers
- fork
- work_directory
- template_directory

Populate each work directory created by Dakota with the contents of a template directory

Specification

Alias: none
Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		copy	Copy template files & directories instead of linking
	Optional		replace	Overwrite existing files within a work directory

Description

Upon creation of a working directory, all the files in the directory specified by `template_directory` are linked (file system symbolic links) or copied into the working directory.

copy

- Keywords Area
- interface
- analysis_drivers
- fork
- work_directory
- template_directory
- copy

Copy template files & directories instead of linking

Specification

Alias: none
Argument(s): none

Description

By default, files & directories specified by the `template_files` and `template_directory` keywords are linked into each work directory created by Dakota.

If the `copy` keyword is specified, these files & directories are copied rather than linked.

replace

- Keywords Area
- interface
- analysis_drivers
- fork
- work_directory
- template_directory
- replace

Overwrite existing files within a work directory

Specification

Alias: none

Argument(s): none

Description

By default, Dakota will not overwrite any existing files in a work directory it creates. The `replace` keyword changes this behavior to force overwriting.

template_files

- Keywords Area
- interface
- analysis_drivers
- fork
- work_directory
- template_files

Populate each work directory created by Dakota with the specified files

Specification

Alias: none

Argument(s): STRINGLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		copy	Copy template files & directories instead of linking
	Optional		replace	Overwrite existing files within a work directory

Description

Upon creation of a working directory, all the files named after the `template_files` keyword are linked (file system symbolic links) into the working directory.

copy

- Keywords Area
- interface
- analysis_drivers
- fork
- work_directory
- template_files
- copy

Copy template files & directories instead of linking

Specification

Alias: none
Argument(s): none

Description

By default, files & directories specified by the `template_files` and `template_directory` keywords are linked into each work directory created by Dakota.

If the `copy` keyword is specified, these files & directories are copied rather than linked.

replace

- Keywords Area
- interface
- analysis_drivers
- fork
- work_directory

- template_files
- replace

Overwrite existing files within a work directory

Specification

Alias: none
Argument(s): none

Description

By default, Dakota will not overwrite any existing files in a work directory it creates. The `replace` keyword changes this behavior to force overwriting.

direct

- Keywords Area
- interface
- analysis_drivers
- direct

Run analysis drivers that are linked-to or compiled-with Dakota

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		processors_per_-analysis	Specify the number of processors per analysis when Dakota is run in parallel

Description

The primary use of the direct interface is to invoke internal test functions that perform parameter to response mappings for simple functions as inexpensively as possible. These problems are compiled directly into the Dakota executable as part of the direct function interface class and are used for algorithm testing.

Dakota also supports direct interfaces to a few select simulation codes. One example is ModelCenter, a commercial simulation management framework from Phoenix Integration. To utilize this interface, a user must first define the simulation specifics within a ModelCenter session and then save these definitions to a ModelCenter configuration file. The `analysis_components` specification provides the means to communicate this configuration file to Dakota's ModelCenter interface.

Examples

The rosenbrock function is available as an executable, which can be launched with fork, and is also compiled with Dakota. The internal version can be launched with:

```
interface
  analysis_drivers = 'rosenbrock'
    direct
```

processors_per_analysis

- Keywords Area
- interface
- analysis_drivers
- direct
- processors_per_analysis

Specify the number of processors per analysis when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

For direct function interfaces, `processors_per_analysis` is an additional optional setting within the required group which can be used to specify multiprocessor analysis partitions. As with the `evaluation_servers`, `analysis_servers`, `evaluation_self_scheduling`, `evaluation_static_scheduling`, `analysis_self_scheduling`, and `analysis_static_scheduling` specifications, `processors_per_analysis` provides a means for the user to override the automatic parallel configuration (refer to Parallel-Library and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"]) for the number of processors used for each analysis partition. Note that if both `analysis_servers` and `processors_per_analysis` are specified and they are not in agreement, then `analysis_servers` takes precedence.

matlab

- Keywords Area
- interface
- analysis_drivers
- matlab

Run Matlab with a direct interface - requires special Dakota build

Specification

Alias: none
Argument(s): none

Description

Dakota supports library-linked interfaces to Matlab, Scilab, and Python scientific computation software, but they must be explicitly enabled when compiling Dakota from source. First consult the Users Manual [[4] "Adams et al., 2010"] for discussion and examples.

Contact the Dakota users mailing list for assistance building and using Dakota with these interfaces.

In all these interfaces, the `analysis_driver` is used to specify a Matlab, Scilab, or Python file which implements the parameter to response mapping.

python

- Keywords Area
- interface
- analysis_drivers
- python

Run Python with a direct interface - requires special Dakota build

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		numpy	Enable the use of numpy in Dakota's Python interface

Description

Dakota supports library-linked interfaces to Matlab, Scilab, and Python scientific computation software, but they must be explicitly enabled when compiling Dakota from source. First consult the Users Manual [[4] "Adams et al., 2010"] for discussion and examples.

Contact the Dakota users mailing list for assistance building and using Dakota with these interfaces.

In all these interfaces, the `analysis_driver` is used to specify a Matlab, Scilab, or Python file which implements the parameter to response mapping.

numpy

- Keywords Area
- interface
- analysis_drivers

- python
- numpy

Enable the use of numpy in Dakota's Python interface

Specification

Alias: none
Argument(s): none

Description

When the `numpy` keyword is used, Dakota expects responses in the form of a Python dictionary of numpy arrays. See the example in `examples/linked_interfaces/Python`.

scilab

- Keywords Area
- interface
- analysis_drivers
- scilab

Run Scilab with a direct interface - requires special Dakota build

Specification

Alias: none
Argument(s): none

Description

Dakota supports library-linked interfaces to Matlab, Scilab, and Python scientific computation software, but they must be explicitly enabled when compiling Dakota from source. First consult the Users Manual [[4] "Adams et al., 2010"] for discussion and examples.

Contact the Dakota users mailing list for assistance building and using Dakota with these interfaces.

In all these interfaces, the `analysis_driver` is used to specify a Matlab, Scilab, or Python file which implements the parameter to response mapping.

grid

- Keywords Area
- interface
- analysis_drivers
- grid

Experimental capability

Specification

Alias: none
Argument(s): none

Description

For grid interfaces, no additional specifications are used at this time.

This capability has been used for interfaces with IDEA and JAVASpaces in the past and is currently a placeholder for future work with Condor and/or Globus services. It is not currently operational.

failure_capture

- Keywords Area
- interface
- analysis_drivers
- failure_capture

Determine how Dakota responds to analysis driver failure

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	abort	(Default) Abort the Dakota job
			retry	Rerun failed analyses
			recover	Substitute dummy values for the responses
			continuation	Cause Dakota to step toward the failed ”target” simulation from a nearby successful ”source”

Description

Dakota can deal with analysis failure in a few ways. However, the first step is that Dakota must detect analysis failure.

- Dakota always expects a results file to be written by the analysis driver
- If the file does not exist when the analysis driver exits, a Dakota error results.
- The analysis driver must communicate an analysis failure to Dakota by writing a results file beginning with ”fail”.
  - This is not case sensitive
  - Everything after fail is ignored

Once Dakota detects analysis failure, the failure can be mitigated in four ways:

- abort (the default)
- retry
- recover
- continuation

Refer to the Simulation Code Failure Capturing chapter of the Users Manual [[4] ”Adams et al., 2010”] for additional information.

abort

- Keywords Area
- interface
- analysis_drivers
- failure_capture
- abort

(Default) Abort the Dakota job

Specification

Alias: none
Argument(s): none

Description

`abort` will stop the Dakota job, as well as any other running analysis drivers.

retry

- Keywords Area
- interface
- analysis_drivers
- failure_capture
- retry

Rerun failed analyses

Specification

Alias: none
Argument(s): INTEGER

Description

The `retry` selection supports an integer input for specifying a limit on retries

recover

- Keywords Area
- interface
- analysis_drivers
- failure_capture
- recover

Substitute dummy values for the responses

Specification

Alias: none
Argument(s): REALLIST

Description

The `recover` selection supports a list of reals for specifying the dummy function values (only zeroth order information is supported) to use for the failed function evaluation.

continuation

- Keywords Area
- interface
- analysis_drivers
- failure_capture
- continuation

Cause Dakota to step toward the failed ”target” simulation from a nearby successful ”source”

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): none

Description

When `failure_capture continuation` is enabled and an evaluation fails, then Dakota will attempt to march incrementally from a previous good evaluation (the "source") toward the failing one (the "target"). Further details about the algorithm employed by Dakota are supplied in the User's Manual [[4] "Adams et al., 2010"].

deactivate

- Keywords Area
- interface
- analysis_drivers
- deactivate

Deactivate Dakota features to simplify interface development, increase execution speed, or reduce memory and disk requirements

Topics

This keyword is related to the topics:

- problem

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		active_set_vector	Deactivate the Active Set Vector
	Optional		evaluation_cache	Do not retain function evaluation history in memory

	Optional	strict_cache_-equality	Do not require strict cache equality when finding duplicates
	Optional	restart_file	Deactivate writing to the restart file

Description

The optional `deactivate` specification block includes three features which a user may deactivate in order to simplify interface development, increase execution speed, and/or reduce memory and disk requirements:

- Active set vector (ASV) control: deactivation of this feature using a `deactivate active_set_vector` specification allows the user to turn off any variability in ASV values so that active set logic can be omitted in the user's simulation interface. This option trades some efficiency for simplicity in interface development. The default behavior is to request the minimum amount of data required by an algorithm at any given time, which implies that the ASV values may vary from one function evaluation to the next. Since the user's interface must return the data set requested by the ASV values, this interface must contain additional logic to account for any variations in ASV content. Deactivating this ASV control causes Dakota to always request a "full" data set (the full function, gradient, and Hessian data that is available from the interface as specified in the responses specification) on each function evaluation. For example, if ASV control has been deactivated and the responses section specifies four response functions, analytic gradients, and no Hessians, then the ASV on every function evaluation will be { 3 3 3 3 }, regardless of what subset of this data is currently needed. While wasteful of computations in many instances, this simplifies the interface and allows the user to return the same data set on every evaluation. Conversely, if ASV control is active (the default behavior), then the ASV requests in this example might vary from { 1 1 1 1 } to { 2 0 0 2 }, etc., according to the specific data needed on a particular function evaluation. This will require the user's interface to read the ASV requests and perform the appropriate logic in conditionally returning only the data requested. In general, the default ASV behavior is recommended for the sake of computational efficiency, unless interface development time is a critical concern. Note that in both cases, the data returned to Dakota from the user's interface must match the ASV passed in, or else a response recovery error will result. However, when the ASV control is deactivated, the ASV values are invariant and need not be checked on every evaluation. *Note:* Deactivating the ASV control can have a positive effect on load balancing for parallel Dakota executions. Thus, there is significant overlap in this ASV control option with speculative gradients. There is also overlap with the mode override approach used with certain optimizers to combine individual value, gradient, and Hessian requests.

- Function evaluation cache: deactivation of this feature using a `deactivate evaluation_cache` specification allows the user to avoid retention of the complete function evaluation history in memory. This can be important for reducing memory requirements in large-scale applications (i.e., applications with a large number of variables or response functions) and for eliminating the overhead of searching for duplicates within the function evaluation cache prior to each new function evaluation (e.g., for improving speed in problems with 1000's of inexpensive function evaluations or for eliminating overhead when performing timing studies). However, the downside is that unnecessary computations may be performed since duplication in function evaluation requests may not be detected. For this reason, this option is not recommended when function evaluations are costly. *Note:* duplication detection within Dakota can be deactivated, but duplication detection features within specific optimizers may still be active.

- Strict Cache Equality: By default, Dakota's evaluation cache and restart capabilities are based on strict binary equality. This provides a performance advantage, as it permits a hash-based data structure to be used

to search the evaluation cache. However, deactiving strict equality may prevent cache misses, which can occur when attempting to use a restart file on a machine different from the one on which it was generated.

- Restart file: deactivation of this feature using a `deactivate restart_file` specification allows the user to eliminate the output of each new function evaluation to the binary restart file. This can increase speed and reduce disk storage requirements, but at the expense of a loss in the ability to recover and continue a run that terminates prematurely (e.g., due to a system crash or network problem). This option is not recommended when function evaluations are costly or prone to failure. Please note that using the `deactivate restart_file` specification will result in a zero length restart file with the default name `dakota.rst`.

These three features may be deactivated independently and concurrently.

active_set_vector

- Keywords Area
- interface
- analysis_drivers
- deactivate
- active_set_vector

Deactivate the Active Set Vector

Specification

Alias: none
Argument(s): none

Description

Described on parent page

evaluation_cache

- Keywords Area
- interface
- analysis_drivers
- deactivate
- evaluation_cache

Do not retain function evaluation history in memory

Specification

Alias: none
Argument(s): none

Description

Described on parent page

strict_cache_equality

- Keywords Area
- interface
- analysis_drivers
- deactivate
- strict_cache_equality

Do not require strict cache equality when finding duplicates

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		cache_tolerance	Specify tolerance when identifying duplicate function evaluations

Description

Described on parent page

cache_tolerance

- Keywords Area
- interface
- analysis_drivers
- deactivate
- strict_cache_equality
- cache_tolerance

Specify tolerance when identifying duplicate function evaluations

Specification

Alias: none

Argument(s): REAL

Description

Described on parent page

restart_file

- Keywords Area
- interface
- analysis_drivers
- deactivate
- restart_file

Deactivate writing to the restart file

Specification

Alias: none
Argument(s): none

Description

Described on parent page

5.5.4 asynchronous

- Keywords Area
- interface
- asynchronous

Specify analysis driver concurrency, when Dakota is run in serial

Topics

This keyword is related to the topics:

- problem
- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Optional	evaluation_-concurrency	Determine how many concurrent evaluations Dakota will schedule
	Optional	local_evaluation_-scheduling	Control how local asynchronous jobs are scheduled
	Optional	analysis_-concurrency	Limit the number of analysis drivers within an evaluation that Dakota will schedule

Description

The optional `asynchronous` keyword specifies use of asynchronous protocols (i.e., background system calls, nonblocking forks, POSIX threads) when evaluations or analyses are invoked. The `evaluation_concurrency` and `analysis_concurrency` specifications serve a dual purpose:

- when running Dakota on a single processor in `asynchronous` mode, the default concurrency of evaluations and analyses is all concurrency that is available. The `evaluation_concurrency` and `analysis_concurrency` specifications can be used to limit this concurrency in order to avoid machine overload or usage policy violation.
- when running Dakota on multiple processors in message passing mode, the default concurrency of evaluations and analyses on each of the servers is one (i.e., the parallelism is exclusively that of the message passing). With the `evaluation_concurrency` and `analysis_concurrency` specifications, a hybrid parallelism can be selected through combination of message passing parallelism with asynchronous parallelism on each server.

evaluation_concurrency

- Keywords Area
- interface
- asynchronous
- evaluation_concurrency

Determine how many concurrent evaluations Dakota will schedule

Topics

This keyword is related to the topics:

- problem
- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

When `asynchronous` execution is enabled, the default behavior is to launch all available evaluations simultaneously. The `evaluation_concurrency` keywrod can be used to limit the number of concurrent evaluations.

local_evaluation_scheduling

- Keywords Area
- interface
- asynchronous
- local_evaluation_scheduling

Control how local asynchronous jobs are scheduled

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	dynamic	Dynamic local scheduling (sequential)
			static	Static local scheduling (tiled)

Description

When performing asynchronous local evaluations, the local_evaluation_scheduling keyword controls how new evaluation jobs are dispatched when one completes.

The two options are:

- `dynamic`
- `static`

If the local_evaluation_scheduling is specified as dynamic (the default), each completed evaluation will be replaced by the next in the local evaluation queue.

If local_evaluation_scheduling is specified as static, each completed evaluation will be replaced by an evaluation number that is congruent modulo the evaluation_concurrency. This is helpful for relative node scheduling as described in Dakota/examples/parallelism. For example, assuming only asynchronous local concurrency (no MPI), if the local concurrency is 6 and job 2 completes, it will be replaced with job 8.

For the case of hybrid parallelism, static local scheduling results in evaluation replacements that are modulo the total capacity, defined as the product of the evaluation concurrency and the number of evaluation servers. Both of these cases can result in idle processors if runtimes are non-uniform, so the default dynamic scheduling is preferred when relative node scheduling is not required.

dynamic

- Keywords Area
- interface
- asynchronous
- local_evaluation_scheduling
- dynamic

Dynamic local scheduling (sequential)

Specification

Alias: none
Argument(s): none

Description

If the local_evaluation_scheduling is specified as dynamic (the default), each completed evaluation will be replaced by the next in the local evaluation queue.

static

- Keywords Area
- interface
- asynchronous
- local_evaluation_scheduling
- static

Static local scheduling (tiled)

Specification

Alias: none
Argument(s): none

Description

If local_evaluation_scheduling is specified as static, each completed evaluation will be replaced by an evaluation number that is congruent modulo the evaluation_concurrency. This is helpful for relative node scheduling as described in Dakota/examples/parallelism. For example, assuming only asynchronous local concurrency (no MPI), if the local concurrency is 6 and job 2 completes, it will be replaced with job 8.

For the case of hybrid parallelism, static local scheduling results in evaluation replacements that are modulo the total capacity, defined as the product of the evaluation concurrency and the number of evaluation servers. Both of these cases can result in idle processors if runtimes are non-uniform, so the default dynamic scheduling is preferred when relative node scheduling is not required.

analysis_concurrency

- Keywords Area
- interface
- asynchronous
- analysis_concurrency

Limit the number of analysis drivers within an evaluation that Dakota will schedule

Topics

This keyword is related to the topics:

- problem
- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

When `asynchronous` execution is enabled and each evaluation involves multiple analysis drivers, then the default behavior is to launch all drivers simultaneously. The `analysis_concurrency` keyword can be used to limit the number of concurrently run drivers.

5.5.5 evaluation_servers

- Keywords Area
- interface
- evaluation_servers

Specify the number of evaluation servers when Dakota is run in parallel

Topics

This keyword is related to the topics:

- problem
- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `evaluation_servers` specification supports user override of the automatic parallel configuration for the number of evaluation servers. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired number of partitions at the evaluation parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

5.5.6 evaluation_scheduling

- Keywords Area
- interface
- evaluation_scheduling

Specify the scheduling of concurrent evaluations when Dakota is run in parallel

Specification

Alias: none
Argument(s): none

	Required/-Optional **Required***(Choose One)*	**Description of Group** **Group 1**	**Dakota Keyword** master	**Dakota Keyword Description** Specify a dedicated master partition for parallel evaluation scheduling
			peer	Specify a peer partition for parallel evaluation scheduling

Description

When Dakota is run in parallel, the partition type and scheduling for the evaluation servers are determined automatically. If these settings are undesirable, they may be overridden by the user using the `evaluation_scheduling` keyword.

The partition type and scheduling are

master

- Keywords Area
- interface
- evaluation_scheduling
- master

Specify a dedicated master partition for parallel evaluation scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a dedicated master partition. In a dedicated master partition, one processor (the "master") dynamically schedules work on the evaluation servers. This reduces the number of processors available to create servers by 1.

peer

- Keywords Area
- interface
- evaluation_scheduling
- peer

Specify a peer partition for parallel evaluation scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required(*Choose One*)	**Group 1**	dynamic	Specify dynamic scheduling in a peer partition when Dakota is run in parallel.
			static	Specify static scheduling in a peer partition when Dakota is run in parallel.

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a peer partition. In a peer partition, all processors are available to be assigned to evaluation servers. The scheduling,`static` or `dynamic`, must also be specified.

dynamic

- Keywords Area
- interface
- evaluation_scheduling
- peer
- dynamic

Specify dynamic scheduling in a peer partition when Dakota is run in parallel.

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

In `dynamic` scheduling, evaluations are assigned to servers as earlier evaluations complete. Dynamic scheduling is advantageous when evaluations are of uneven duration.

static

- Keywords Area
- interface
- evaluation_scheduling
- peer
- static

Specify static scheduling in a peer partition when Dakota is run in parallel.

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

In `static` scheduling, all available evaluations are assigned to servers in a predetermined fashion. Each completed evaluation is replaced with one congruent modulo the evaluation concurrency. For example, with 6 servers, eval number 2 will be replaced by eval number 8.

5.5.7 processors_per_evaluation

- Keywords Area
- interface
- processors_per_evaluation

Specify the number of processors per evaluation server when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `processors_per_evaluation` specification supports user override of the automatic parallel configuration for the number of processors in each evaluation server. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired server size at the evaluation parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

5.5.8 analysis_servers

- Keywords Area
- interface
- analysis_servers

Specify the number of analysis servers when Dakota is run in parallel

Topics

This keyword is related to the topics:

- problem
- concurrency_and_parallelism

Specification

Alias: none
Argument(s): INTEGER

Description

The optional `analysis_servers` specification supports user override of the automatic parallel configuration for the number of analysis servers. That is, if the automatic configuration is undesirable for some reason, the user can enforce a desired number of partitions at the analysis parallelism level. Refer to ParallelLibrary and the Parallel Computing chapter of the Users Manual [[4] "Adams et al., 2010"] for additional information.

5.5.9 analysis_scheduling

- Keywords Area
- interface
- analysis_scheduling

Specify the scheduling of concurrent analyses when Dakota is run in parallel

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	master	Specify a dedicated master partition for parallel analysis scheduling
			peer	Specify a peer partition for parallel analysis scheduling

Description

When Dakota is run in parallel, the partition type for the analysis servers is determined automatically. If this setting is undesirable, it may be overridden by the user using the `analysis_scheduling` keyword.

The partition type and scheduling are

master

- Keywords Area
- interface
- analysis_scheduling
- master

Specify a dedicated master partition for parallel analysis scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none

Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a dedicated master partition. In a dedicated master partition, one processor (the ”master”) dynamically schedules work on the analysis servers. This reduces the number of processors available to create servers by 1.

peer

- Keywords Area
- interface
- analysis_scheduling
- peer

Specify a peer partition for parallel analysis scheduling

Topics

This keyword is related to the topics:

- concurrency_and_parallelism

Specification

Alias: none
Argument(s): none

Description

This option overrides the Dakota parallel automatic configuration, forcing the use of a peer partition. In a peer partition, all processors are available to be assigned to analysis servers. Note that unlike the case of `evaluation_scheduling`, it is not possible to specify `static` or `dynamic`.

5.6 responses

- Keywords Area
- responses

Description of the model output data returned to Dakota upon evaluation of an interface.

Topics

This keyword is related to the topics:

- block

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Optional		id_responses	Name the response block, when there are multiple responses blocks
	Optional		descriptors	Labels for the responses
	Required*(Choose One)*	**Group 1**	objective_functions	Response type suitable for optimization
			calibration_terms	Response type suitable for calibration or least squares
			response_functions	Generic response type
	Required*(Choose One)*	**Group 2**	no_gradients	Gradients will not be used
			analytic_gradients	Analysis driver will return gradients
			mixed_gradients	Gradients are needed and will be obtained from a mix of numerical and analytic sources
			numerical_-gradients	Gradients are needed and will be approximated by finite differences
	Required*(Choose One)*	**Group 3**	no_hessians	Hessians will not be used
			numerical_hessians	Hessians are needed and will be approximated by finite differences
			quasi_hessians	Hessians are needed and will be approximated by secant updates (BFGS or SR1) from a series of gradient evaluations

			analytic_hessians	Hessians are needed and are available directly from the analysis driver
			mixed_hessians	Hessians are needed and will be obtained from a mix of numerical, analytic, and ”quasi” sources

Description

The `responses` specification in a Dakota input file indicates the types of data that can be returned by an interface when invoked during Dakota's execution. The specification includes three groups and two optional keywords.

Group 1 is related to the type and number of responses expected by Dakota

The specification must be one of three types:

1. objective and constraint functions
2. calibration (least squares) terms and constraint functions
3. a generic response functions specification.

These correspond to optimization, least squares, and uncertainty quantification methods, respectively. The response type chosen from Group 1 should be consistent with the iterative technique called for in the method specification. Certain general-purpose iterative techniques, such as parameter studies and design of experiments methods, can be used with any of these data sets.

Each type of response has additional required and optional keywords.

Group 2 is related to the availability of first derivatives (gradient vectors) for the response functions.

The gradient specification also links back to the iterative method used. Gradients commonly are needed when the iterative study involves gradient-based optimization, local reliability analysis for uncertainty quantification, or local sensitivity analysis. They can optionally be used to build some types of surrogate models.

Group 3 is related to the availability of second derivatives (Hessian matrices) for the response functions.

Hessian availability for the response functions is similar to the gradient availability specifications, with the addition of support for ”quasi-Hessians”. The Hessian specification also links back to the iterative method in use; Hessians commonly would be used in gradient-based optimization by full Newton methods or in reliability analysis with second-order limit state approximations or second-order probability integrations.

Examples

Several examples follow. The first example shows an optimization data set containing an objective function and two nonlinear inequality constraints. These three functions have analytic gradient availability and no Hessian availability.

```
responses
  objective_functions = 1
    nonlinear_inequality_constraints = 2
  analytic_gradients
  no_hessians
```

The next example shows a typical specification for a calibration data set. The six residual functions will have numerical gradients computed using the dakota finite differencing routine with central differences of 0.1% (plus/minus delta relative to current variables value = .001∗value).

```
responses
  calibration_terms = 6
  numerical_gradients
    method_source dakota
    interval_type central
    fd_gradient_step_size = .001
  no_hessians
```

The last example shows a generic specification that could be used with a nondeterministic sampling iterator. The three response functions have no gradient or Hessian availability; therefore, only function values will be used by the iterator.

```
responses
  response_functions = 3
  no_gradients
  no_hessians
```

Parameter study and design of experiments iterators are not restricted in terms of the response data sets which may be catalogued; they may be used with any of the function specification examples shown above.

Theory

Responses specify the total data set that is available for use by the method over the course of iteration. This is distinguished from the data subset described by an active set vector (see Dakota File Data Formats in the Users Manual [Adams et al., 2010]) indicating the particular subset of the response data needed for a particular function evaluation. Thus, the responses specification is a broad description of the data to be used during a study whereas the active set vector indicates the subset currently needed.

5.6.1 id_responses

- Keywords Area
- responses
- id_responses

Name the response block, when there are multiple responses blocks

Topics

This keyword is related to the topics:

- block_identifier

Specification

Alias: none
Argument(s): STRING

Description

The optional block identifier `id_responses` specifies a string to uniquely identify a particular responses specification (typically used when there are multiple present). A model can then specify or point to this response set by specifying the same string in its `responses_pointer` specification. For example, a model whose specification contains `responses_pointer = 'R1'` will use a responses set with `id_responses = 'R1'`.

If the `id_responses` specification is omitted, a particular responses specification will be used by a model only if that model omits specifying a `responses_pointer` and if the responses set was the last set parsed (or is the only set parsed). In common practice, if only one responses set exists, then `id_responses` can be safely omitted from the responses specification and `responses_pointer` can be omitted from the model specification(s), since there is no potential for ambiguity in this case.

5.6.2 descriptors

- Keywords Area
- responses
- descriptors

Labels for the responses

Specification

Alias: response_descriptors
Argument(s): STRINGLIST

Description

The optional response labels specification `descriptors` is a list of strings which will be printed in Dakota output to identify the values for particular response functions.

Note that the ordering of responses and descriptors in the input **currently** must match the order of the values returned to Dakota in a results_file. See the example below.

The default descriptor strings use a root string plus a numeric identifier. This root string is

- `"obj_fn"` for objective functions
- `"least_sq_term"` for least squares terms
- `"response_fn"` for generic response functions
- `"nln_ineq_con"` for nonlinear inequality constraints
- `"nln_eq_con"` for nonlinear equality constraints

Examples

Note that the descriptors **currently** must match the order of the values in the results file; they are not used to validate the returned data. For example, if the `responses` block contains:

```
\c descriptors 'x1' 'x2' 'x3'
```

and the results file contains

```
4 x1
5 x3
6 x2
```

Then Dakota will understand the returned data to be:

```
x1 = 4
x2 = 5
x3 = 6
```

5.6.3 objective_functions

- Keywords Area
- responses
- objective_functions

Response type suitable for optimization

Specification

Alias: num_objective_functions

Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		sense	Whether to minimize or maximize each objective function
	Optional		primary_scale_-types	Choose a scaling type for each response
	Optional		primary_scales	Supply a characteristic value to scale each reponse
	Optional		weights	Specify weights for each objective function
	Optional		nonlinear_-inequality_-constraints	Group to specify nonlinear inequality constraints

	Optional	nonlinear_equality-_constraints	Group to specify nonlinear equality constraints

Description

The `objective_functions` keyword specifies the number of objective functions returned to Dakota. The number of objective functions must be 1 or greater.

Constraints

The keywords `nonlinear_inequality_constraints`, and `nonlinear_equality_constraints` specify the number of nonlinear inequality constraints, and nonlinear equality constraints, respectively. When interfacing to external applications, the responses must be returned to Dakota in this order: objective functions, nonlinear_inequality_constraints, then nonlinear_equality_constraints.

Any linear constraints present in an application need only be input to an optimizer at start up and do not need to be part of the data returned on every function evaluation. These are therefore specified in the method block.

Bounds on the design variables are specified in the variables block.

Optional Keywords

The optional keywords relate to scaling the objective functions (for better numerical results), formulating the problem as minimization or maximization, and dealing with multiple objective functions. If scaling is used, it is applied before multi-objective weighted sums are formed.

See Also

These keywords may also be of interest:

- calibration_terms
- response_functions
- method
- variables

sense

- Keywords Area
- responses
- objective_functions
- sense

Whether to minimize or maximize each objective function

Specification

Alias: none

Argument(s): STRINGLIST

Description

The `sense` keyword is used to declare whether each objective function should be minimized or maximized. The argument options are:

- `"minimization"`
- `"maximization"` These can be abbreviated to `"min"` and `"max"`.

The number of strings should either be equal to the number of objective functions, or one. If a single string is specified it will apply to each objective function.

primary_scale_types

- Keywords Area
- responses
- objective_functions
- primary_scale_types

Choose a scaling type for each response

Specification

Alias: objective_function_scale_types
Argument(s): STRINGLIST

Description

The `primary_scale_types` keyword specifies one of number of primary functions strings indicating the scaling type for each response value in methods that support scaling, when scaling is enabled.

See the scaling page for details on how to use this keyword. Note that primary response functions (objective, calibration, or response functions) cannot be automaticaly scaled due to lack of bounds, so valid scale types are `'none'` `'value'` and `'log'`.

primary_scales

- Keywords Area
- responses
- objective_functions
- primary_scales

Supply a characteristic value to scale each reponse

Specification

Alias: objective_function_scales
Argument(s): REALLIST

Description

Each entry in `primary_scales` is a user-specified nonzero characteristic value to scale each response.

The argument may be of length 1 or the number of primary response functions. See the scaling page for details on how to use this keyword.

weights

- Keywords Area
- responses
- objective_functions
- weights

Specify weights for each objective function

Specification

Alias: multi_objective_weights
Argument(s): REALLIST

Description

If the number of objective functions is greater than 1, then a `weights` specification provides a simple weighted-sum approach to combining multiple objectives into a single objective:

$$f = \sum_{i=1}^{n} w_i f_i$$

If weights are not specified, then each response is given equal weighting:

$$f = \sum_{i=1}^{n} \frac{f_i}{n}$$

where, in both of these cases, a ”minimization” sense will retain a positive weighting for a minimizer and a ”maximization” sense will apply a negative weighting.

nonlinear_inequality_constraints

- Keywords Area
- responses
- objective_functions
- nonlinear_inequality_constraints

Group to specify nonlinear inequality constraints

Specification

Alias: num_nonlinear_inequality_constraints
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		lower_bounds	Specify minimum values
	Optional		upper_bounds	Specify maximium values
	Optional		scale_types	Choose how each constraint is scaled
	Optional		scales	Characteristic values for scaling

Description

The `lower_bounds` and `upper_bounds` specifications provide the lower and upper bounds for 2-sided nonlinear inequalities of the form

$$g_l \leq g(x) \leq g_u$$

The defaults for the inequality constraint bounds are selected so that one-sided inequalities of the form

$$g(x) \leq 0.0$$

result when there are no user constraint bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity. This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`). The same approach is used for nonexistent linear inequality bounds and for nonexistent design variable bounds.

The `scale_types` and `scales` keywords are related to scaling of $g(x)$. See the scaling page for details.

lower_bounds

- Keywords Area
- responses
- objective_functions
- nonlinear_inequality_constraints
- lower_bounds

Specify minimum values

Specification

Alias: nonlinear_inequality_lower_bounds
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- responses
- objective_functions
- nonlinear_inequality_constraints
- upper_bounds

Specify maximium values

Specification

Alias: nonlinear_inequality_upper_bounds
Argument(s): REALLIST

Description

Specify maximium values

scale_types

- Keywords Area
- responses
- objective_functions
- nonlinear_inequality_constraints
- scale_types

Choose how each constraint is scaled

Specification

Alias: nonlinear_inequality_scale_types
Argument(s): STRINGLIST

Description

See the scaling page for details on how to use this keyword.

scales

- Keywords Area
- responses
- objective_functions
- nonlinear_inequality_constraints
- scales

Characteristic values for scaling

Specification

Alias: nonlinear_inequality_scales
Argument(s): REALLIST

Description

See the scaling page for details on how to use this keyword.

nonlinear_equality_constraints

- Keywords Area
- responses
- objective_functions
- nonlinear_equality_constraints

Group to specify nonlinear equality constraints

Specification

Alias: num_nonlinear_equality_constraints
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		targets	Target values for the nonlinear equality constraint
	Optional		scale_types	Choose how each constraint is scaled
	Optional		scales	Characteristic values for scaling

Description

The `targets` specification provides the targets for nonlinear equalities of the form

$$g(x) = g_t$$

and the defaults for the equality targets enforce a value of `0.` for each constraint

$$g(x) = 0.0$$

The `scale_types` and `scales` keywords are related to scaling of $g(x)$. See the scaling page for details.

targets

- Keywords Area
- responses
- objective_functions
- nonlinear_equality_constraints
- targets

Target values for the nonlinear equality constraint

Specification

Alias: nonlinear_equality_targets
Argument(s): REALLIST

Description

The `targets` specification provides the targets for nonlinear equalities of the form

$$g(x) = g_t$$

and the defaults for the equality targets enforce a value of `0.0` for each constraint:

$$g(x) = 0.0$$

scale_types

- Keywords Area
- responses
- objective_functions
- nonlinear_equality_constraints
- scale_types

Choose how each constraint is scaled

Specification

Alias: nonlinear_equality_scale_types
Argument(s): STRINGLIST

Description

See the scaling page for details on how to use this keyword.

scales

- Keywords Area
- responses
- objective_functions
- nonlinear_equality_constraints
- scales

Characteristic values for scaling

Specification

Alias: nonlinear_equality_scales
Argument(s): REALLIST

Description

See the scaling page for details on how to use this keyword.

5.6.4 calibration_terms

- Keywords Area
- responses
- calibration_terms

Response type suitable for calibration or least squares

Specification

Alias: least_squares_terms num_least_squares_terms
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		primary_scale_-types	Choose a scaling type for each response
	Optional		primary_scales	Supply a characteristic value to scale each reponse

	Optional	weights	Apply different weights to each response
	Optional	calibration_data_-file	Supply calibration data in a text file
	Optional	nonlinear_-inequality_-constraints	Group to specify nonlinear inequality constraints
	Optional	nonlinear_equality-_constraints	Group to specify nonlinear equality constraints

Description

Responses for a calibration study are specified using `calibration_terms` and optional keywords for weighting/scaling, data, and constraints. In general when calibrating, Dakota automatically tunes parameters θ to minimize discrepancies or residuals between the model and the data:

$$R_i = y_i^{Model}(\theta) - y_i^{Data}.$$

There are two use cases:

- If calibration_data_file is NOT specified, then each of the calibration terms returned to Dakota through the interface is a residual R_i to be driven toward zero.
- If calibration_data_file IS specified, then each of the calibration terms returned to Dakota must be a response $y_i^{Model}(\theta)$, which Dakota will difference with the data in the specified data file.

Constraints

The keywords `nonlinear_inequality_constraints`, and `nonlinear_equality_constraints` specify the number of nonlinear inequality constraints, and nonlinear equality constraints, respectively. When interfacing to external applications, the responses must be returned to Dakota in this order: calibration terms, nonlinear_inequality_constraints, then nonlinear_equality_constraints.

Any linear constraints present in an application need only be input to an optimizer at start up and do not need to be part of the data returned on every function evaluation. These are therefore specified in the method block.

Optional Keywords

The optional keywords relate to scaling responses (for better numerical results), dealing with multiple residuals, and importing data.

See scaling for more details on scaling. If scaling is specified, then it is applied to each residual prior to squaring:

$$f = \sum_{i=1}^{n} w_i \left(\frac{y_i^{Model} - y_i^{Data}}{s_i}\right)^2$$

In the case where experimental data uncertainties are supplied, then the weights are automatically defined to be the inverse of the experimental variance:

$$f = \sum_{i=1}^{n} \frac{1}{\sigma_i^2} \left(\frac{y_i^{Model} - y_i^{Data}}{s_i}\right)^2$$

Theory

Dakota calibration terms are typically used to solve problems of parameter estimation, system identification, and model calibration/inversion. Local least squares calibration problems are most efficiently solved using special-purpose least squares solvers such as Gauss-Newton or Levenberg-Marquardt; however, they may also be solved using any general-purpose optimization algorithm in Dakota. While Dakota can solve these problems with either least squares or optimization algorithms, the response data sets to be returned from the simulator are different when using objective_functions versus calibration_terms.

Least squares calibration involves a set of residual functions, whereas optimization involves a single objective function (sum of the squares of the residuals), i.e.,

$$f = \sum_{i=1}^{n} R_i^2 = \sum_{i=1}^{n} \left(y_i^{Model}(\theta) - y_i^{Data}\right)^2$$

where f is the objective function and the set of R_i are the residual functions, most commonly defined as the difference between a model response and data. Therefore, function values and derivative data in the least squares case involve the values and derivatives of the residual functions, whereas the optimization case involves values and derivatives of the sum of squares objective function. This means that in the least squares calibration case, the user must return each of `n` residuals separately as a separate calibration term. Switching between the two approaches sometimes requires different simulation interfaces capable of returning the different granularity of response data required, although Dakota supports automatic recasting of residuals into a sum of squares for presentation to an optimization method. Typically, the user must compute the difference between the model results and the observations when computing the residuals. However, the user has the option of specifying the observational data (e.g. from physical experiments or other sources) in a file.

See Also

These keywords may also be of interest:

- objective_functions
- response_functions

primary_scale_types

- Keywords Area
- responses
- calibration_terms
- primary_scale_types

Choose a scaling type for each response

Specification

Alias: calibration_term_scale_types least_squares_term_scale_types
Argument(s): STRINGLIST

Description

The `primary_scale_types` keyword specifies one of number of primary functions strings indicating the scaling type for each response value in methods that support scaling, when scaling is enabled.

See the scaling page for details on how to use this keyword. Note that primary response functions (objective, calibration, or response functions) cannot be automaticaly scaled due to lack of bounds, so valid scale types are `'none'` `'value'` and `'log'`.

primary_scales

- Keywords Area
- responses
- calibration_terms
- primary_scales

Supply a characteristic value to scale each reponse

Specification

Alias: calibration_term_scales least_squares_term_scales

Argument(s): REALLIST

Description

Each entry in `primary_scales` is a user-specified nonzero characteristic value to scale each response.

The argument may be of length 1 or the number of primary response functions. See the scaling page for details on how to use this keyword.

weights

- Keywords Area
- responses
- calibration_terms
- weights

Apply different weights to each response

Specification

Alias: calibration_weights least_squares_weights

Argument(s): REALLIST

Description

The `weights` specification provides a means to specify a relative emphasis among the vector of squared residuals through multiplication of these squared residuals by a vector of weights:

$$f = \sum_{i=1}^{n} w_i R_i^2 = \sum_{i=1}^{n} w_i (y_i^M - y_i^O)^2$$

calibration_data_file

- Keywords Area
- responses
- calibration_terms
- calibration_data_file

Supply calibration data in a text file

Topics

This keyword is related to the topics:

- problem

Specification

Alias: least_squares_data_file

Argument(s): STRING

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		num_experiments	Add context to data: number of different experiments
	Optional		num_replicates	Add context to data: number of repeats for each experiment
	Optional*(Choose One)*	**Group 1**	annotated	(Default) Specified file is in the annotated format
			freeform	Specified file is in the freeform format

	Optional	num_config_-variables	Add context to data: number of configuration variables.
	Optional	num_std_deviations	Add context to data: measurement or observation error

Description

`calibration_data_file` specifies a text file containing data which Dakota will attempt to match by calibrating the model variables. The options for it related to either:

- format: whether the data file is in `annotated`, or `freeform` format
- content: where `num_experiments`, `num_replicates`, `num_config_variables`, and `num_std_deviations` indicate which columns appear in the data.

While some components may be omitted, the most complete version of a an annotated calibration data file could include columns corresponding to:

```
index | configuration xvars | y data observations | y data std devs
```

Here each such row in the file corresponds to an experiment or replicate observation of an experment to be compared to the model output.

Simple case

In the simplest case, no data content descriptors are specified, so the data file must contain only the y^{Data} observations which represent a single experimental observation. In this case, the data file should have N_{terms} columns and 1 row, where N_{terms} is the value of calibration_terms.

For each function evaluation, Dakota will run the analysis driver, which must return N_{terms} model responses. Then the residuals are computed as:

$$R_i = y_i^{Model} - y_i^{Data}.$$

These residuals can be weighted using weights.

With experimental standard deviations

In this case, the keyword `num_std_deviations` is added, where `num_std_deviations` must be set to 1 or N_{terms} , where N_{terms} is the value of calibration_terms.

If set to 1, each row of the datafile will now have N_{terms} y data values followed by 1 experimental standard deviation that will be used for all N_{terms} responses. If set to N_{terms}, each row must contain N_{terms} y data values followed by N_{terms} y standard deviations.

Dakota will run the analysis driver, which must return N_{terms} responses. Then the residuals are computed as:

$$R_i = \frac{y_i^{Model} - y_i^{Data}}{s_i}$$

for $i = 1 \ldots N_{terms}$. In the case with 1 standard deviation, s_i will be the same for all i.

Fully general case

In the most general case, the content of the data file is described by the arguments of four parameters. The parameters are optional, and defaults are described below.

- `num_experiments` (N_{exp})

  Default: $N_{exp} = N_{terms}$

  This indicates that the data represents multiple experiments, where each experiment might be conducted with different values of configuration variables.

- `num_replicates` (N_{rep})

  Default: $N_{rep} = [1 \times N_{exp}]$ vector of 1's

  This indicates the number of replicates for each experiment. Replicates are all conducted at the same values of configuration varaibles. It is a vector of length N_{exp}.

- `num_config_variables` (N_{cfg})

  This is not yet supported, but will specify the values of experimental conditions at which data were collected.

- `num_std_deviations` (N_{std})

  Default: $N_{std} = 0$

The keyword `num_std_deviations` allows the user to specify different values for the experimental error, i.e., measurement or observation error in the `calibration_data_file`. If `num_std_deviations` = 0, the user does not specify any experimental error terms in the `calibration_data_file`, only the actual observations are specified. If the user specifies `num_std_deviations` equal to the number of calibration terms, then each row of a freeform `calibration_data_file` must contain two times `calibration_terms`. The first `calibration_terms` columns are the experimental data, and the second `calibration_terms` columns are the experimental standard deviations. For example, if the user has three `calibration` terms, and specifies `num_std_deviations` = 3, then the calibration data must contain six columns. The first three columns will contain the data, and the second three columns will contain the experimental error for the data in the first three columns. Finally, if the user specifies `num_std_deviations` = 1, the same value of the standard deviations will be used for all of the calibration terms. In the example given above, with three calibration terms, if `num_std_deviations` = 1, the there would be four columns in the `calibration_data_file`, and the fourth column would contain the standard deviation that would be applied to all three calibration columns. Note that the standard deviations are given in units of a standard deviation or sigma term, not a variance term. These standard deviations are used to weight the residuals in the sum-of-squares objective.

A more advanced use of the `calibration_data_file` might specify `num_experiments` N_E and `num_replicates` N_R, indicating that there are multiple experiments, each with 1 or more replicates to inform each residual term. When multiple experiments or replicates are present, Dakota will expand the number of residuals for the repeat measurement data and difference with the data accordingly. For example, if the user has one experiment but five replicates were taken, in the example above with three calibration terms, the `calibration_data_file` would need to contain five rows (one for each replicate), and each row should contain three experimental data values that will be differenced with respect to the appropriate model response. In this example, $N_E = 1$, and $N_R = 5$. One can have varying numbers of experiments and replicates, such as two experiments, one with 3 replicates and one with 7 replicates. In this example, `num_experiments` = 2, and the number of replicates would be given as an integer vector: `num_replicates` = 3 7. In this example, there would be a total of 10 rows of data in the experimental data file. To summarize, Dakota will calculate the sum of the squared residuals as:

$$f = \sum_{i=1}^{N_E} \sum_{j=1}^{N_{R_i}} R_{i,j}^2$$

where the residuals now are calculated as:

$$R_{i,j} = y_i^{Model}(\theta) - y_{i,j}^{Data}.$$

num_experiments

- Keywords Area
- responses
- calibration_terms
- calibration_data_file
- num_experiments

Add context to data: number of different experiments

Specification

Alias: none
Argument(s): INTEGER

Description

See calibration_data_file.

num_replicates

- Keywords Area
- responses
- calibration_terms
- calibration_data_file
- num_replicates

Add context to data: number of repeats for each experiment

Specification

Alias: none
Argument(s): INTEGERLIST

Description

See calibration_data_file.

annotated

- Keywords Area
- responses
- calibration_terms
- calibration_data_file
- annotated

(Default) Specified file is in the annotated format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Annotated Matrix (default for all I/O; specified via `annotated`): text file with one leading row of comments/column labels and one leading column of evaluation/row IDs surrounding num_rows x num_cols whitespace-separated numeric data, (newlines separating rows are not currently required, but may be in the future). The numeric data in a row may correspond to variables, variables followed by responses, data point for calibration, etc., depending on context.

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

freeform

- Keywords Area
- responses
- calibration_terms
- calibration_data_file
- freeform

Specified file is in the freeform format

Topics

This keyword is related to the topics:

- file_formats

Specification

Alias: none
Argument(s): none

Description

Free-form Matrix: text file with no leading row and no leading column. The num_rows x num_cols total numeric data entries may appear separated with any whitespace including arbitrary spaces, tabs, and newlines. In this format, vectors may therefore appear as a single row or single column (or mixture; entries will populate the vector in order).

Important Notes

- Prior to October 2011, calibration and surrogate data files were free-form format. They now default to annotated format, though `freeform` remains an option.
- For both formats, a warning will be generated if a specific number of data are expected, but extra is found and an error generated when there is insufficient data.
- Some TPLs like SCOLIB and JEGA manage their own file I/O and only support the free-form option.

num_config_variables

- Keywords Area
- responses
- calibration_terms
- calibration_data_file
- num_config_variables

Add context to data: number of configuration variables.

Specification

Alias: none

Argument(s): INTEGER

Description

See calibration_data_file.

num_std_deviations

- Keywords Area
- responses
- calibration_terms
- calibration_data_file
- num_std_deviations

Add context to data: measurement or observation error

Specification

Alias: none
Argument(s): INTEGER

Description

See calibration_data_file.

nonlinear_inequality_constraints

- Keywords Area
- responses
- calibration_terms
- nonlinear_inequality_constraints

Group to specify nonlinear inequality constraints

Specification

Alias: num_nonlinear_inequality_constraints
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		lower_bounds	Specify minimum values
	Optional		upper_bounds	Specify maximium values
	Optional		scale_types	Choose how each constraint is scaled
	Optional		scales	Characteristic values for scaling

Description

The `lower_bounds` and `upper_bounds` specifications provide the lower and upper bounds for 2-sided nonlinear inequalities of the form

$$g_l \leq g(x) \leq g_u$$

The defaults for the inequality constraint bounds are selected so that one-sided inequalities of the form

$$g(x) \leq 0.0$$

result when there are no user constraint bounds specifications (this provides backwards compatibility with previous Dakota versions).

In a user bounds specification, any upper bound values greater than `+bigRealBoundSize` (1.e+30, as defined in Minimizer) are treated as +infinity and any lower bound values less than `-bigRealBoundSize` are treated as -infinity. This feature is commonly used to drop one of the bounds in order to specify a 1-sided constraint (just as the default lower bounds drop out since `-DBL_MAX` < `-bigRealBoundSize`). The same approach is used for nonexistent linear inequality bounds and for nonexistent design variable bounds.

The `scale_types` and `scales` keywords are related to scaling of $g(x)$. See the scaling page for details.

lower_bounds

- Keywords Area
- responses
- calibration_terms
- nonlinear_inequality_constraints
- lower_bounds

Specify minimum values

Specification

Alias: nonlinear_inequality_lower_bounds
Argument(s): REALLIST

Description

Specify minimum values

upper_bounds

- Keywords Area
- responses
- calibration_terms
- nonlinear_inequality_constraints
- upper_bounds

Specify maximium values

Specification

Alias: nonlinear_inequality_upper_bounds
Argument(s): REALLIST

Description

Specify maximium values

scale_types

- Keywords Area
- responses
- calibration_terms
- nonlinear_inequality_constraints
- scale_types

Choose how each constraint is scaled

Specification

Alias: nonlinear_inequality_scale_types
Argument(s): STRINGLIST

Description

See the scaling page for details on how to use this keyword.

scales

- Keywords Area
- responses
- calibration_terms
- nonlinear_inequality_constraints
- scales

Characteristic values for scaling

Specification

Alias: nonlinear_inequality_scales
Argument(s): REALLIST

Description

See the scaling page for details on how to use this keyword.

nonlinear_equality_constraints

- Keywords Area
- responses
- calibration_terms
- nonlinear_equality_constraints

Group to specify nonlinear equality constraints

Specification

Alias: num_nonlinear_equality_constraints
Argument(s): INTEGER

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description

	Optional	targets	Target values for the nonlinear equality constraint
	Optional	scale_types	Choose how each constraint is scaled
	Optional	scales	Characteristic values for scaling

Description

The `targets` specification provides the targets for nonlinear equalities of the form

$$g(x) = g_t$$

and the defaults for the equality targets enforce a value of `0.` for each constraint

$$g(x) = 0.0$$

The `scale_types` and `scales` keywords are related to scaling of $g(x)$. See the scaling page for details.

targets

- Keywords Area
- responses
- calibration_terms
- nonlinear_equality_constraints
- targets

Target values for the nonlinear equality constraint

Specification

Alias: nonlinear_equality_targets
Argument(s): REALLIST

Description

The `targets` specification provides the targets for nonlinear equalities of the form

$$g(x) = g_t$$

and the defaults for the equality targets enforce a value of `0.0` for each constraint:

$$g(x) = 0.0$$

scale_types

- Keywords Area
- responses
- calibration_terms
- nonlinear_equality_constraints
- scale_types

Choose how each constraint is scaled

Specification

Alias: nonlinear_equality_scale_types
Argument(s): STRINGLIST

Description

See the scaling page for details on how to use this keyword.

scales

- Keywords Area
- responses
- calibration_terms
- nonlinear_equality_constraints
- scales

Characteristic values for scaling

Specification

Alias: nonlinear_equality_scales
Argument(s): REALLIST

Description

See the scaling page for details on how to use this keyword.

5.6.5 response_functions

- Keywords Area
- responses
- response_functions

Generic response type

Specification

Alias: num_response_functions
Argument(s): INTEGER

Description

A generic response data set is specified using `response_functions`. Each of these functions is simply a response quantity of interest with no special interpretation taken by the method in use.

Whereas objective, constraint, and residual functions have special meanings for optimization and least squares algorithms, the generic response function data set need not have a specific interpretation and the user is free to define whatever functional form is convenient.

Theory

This type of data set is used by uncertainty quantification methods, in which the effect of parameter uncertainty on response functions is quantified, and can also be used in parameter study and design of experiments methods (although these methods are not restricted to this data set), in which the effect of parameter variations on response functions is evaluated.

See Also

These keywords may also be of interest:

- objective_functions
- calibration_terms

5.6.6 no_gradients

- Keywords Area
- responses
- no_gradients

Gradients will not be used

Specification

Alias: none
Argument(s): none

Description

The `no_gradients` specification means that gradient information is not needed in the study. Therefore, it will neither be retrieved from the simulation nor computed with finite differences. The `no_gradients` keyword is a complete specification for this case.

See Also

These keywords may also be of interest:

- numerical_gradients
- analytic_gradients
- mixed_gradients

5.6.7 analytic_gradients

- Keywords Area
- responses
- analytic_gradients

Analysis driver will return gradients

Specification

Alias: none
Argument(s): none

Description

The `analytic_gradients` specification means that gradient information is available directly from the simulation (finite differencing is not required). The simulation must return the gradient data in the Dakota format (enclosed in single brackets; see Dakota File Data Formats in the Users Manual [[4] "Adams et al., 2010"]) for the case of file transfer of data. The `analytic_gradients` keyword is a complete specification for this case.

See Also

These keywords may also be of interest:

- numerical_gradients
- no_gradients
- mixed_gradients

5.6.8 mixed_gradients

- Keywords Area
- responses
- mixed_gradients

Gradients are needed and will be obtained from a mix of numerical and analytic sources

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required		id_numerical_-gradients	Identify which numerical gradient corresponds to which response
	Required		id_analytic_-gradients	Identify which analytical gradient corresponds to which response
	Optional		method_source	Specify which finite difference routine is used
	Optional*(Choose One)*	**Group 1**	dakota	(Default) Use internal Dakota finite differences algorithm
			vendor	Use non-Dakota fd algorithm
	Optional		interval_type	Specify how to compute gradients and hessians
	Optional*(Choose One)*	**Group 2**	forward	Use forward differences
			central	Use central differences
	Optional		fd_step_size	Step size used when computing gradients and Hessians

Description

The `mixed_gradients` specification means that some gradient information is available directly from the simulation (analytic) whereas the rest will have to be finite differenced (numerical). This specification allows the user to make use of as much analytic gradient information as is available and then finite difference for the rest.

The `method_source`, `interval_type`, and `fd_gradient_step_size` specifications pertain to those functions listed by the `id_numerical_gradients` list.

Examples

For example, the objective function may be a simple analytic function of the design variables (e.g., weight) whereas the constraints are nonlinear implicit functions of complex analyses (e.g., maximum stress).

See Also

These keywords may also be of interest:

- numerical_gradients
- no_gradients

- analytic_gradients

id_numerical_gradients

- Keywords Area
- responses
- mixed_gradients
- id_numerical_gradients

Identify which numerical gradient corresponds to which response

Topics

This keyword is related to the topics:

- objective_function_pointer

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `id_analytic_gradients` list specifies by number the functions which have analytic gradients, and the `id_numerical_gradients` list specifies by number the functions which must use numerical gradients. Each function identifier, from 1 through the total number of functions, must appear once and only once within the union of the `id_analytic_gradients` and `id_numerical_gradients` lists.

See Also

These keywords may also be of interest:

- id_analytic_gradients

id_analytic_gradients

- Keywords Area
- responses
- mixed_gradients
- id_analytic_gradients

Identify which analytical gradient corresponds to which response

Topics

This keyword is related to the topics:

- objective_function_pointer

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `id_analytic_gradients` list specifies by number the functions which have analytic gradients, and the `id_numerical_gradients` list specifies by number the functions which must use numerical gradients. Each function identifier, from 1 through the total number of functions, must appear once and only once within the union of the `id_analytic_gradients` and `id_numerical_gradients` lists.

See Also

These keywords may also be of interest:

- id_numerical_gradients

method_source

- Keywords Area
- responses
- mixed_gradients
- method_source

Specify which finite difference routine is used

Specification

Alias: none
Argument(s): none

Description

The `method_source` setting specifies the source of the finite differencing routine that will be used to compute the numerical gradients:

- `dakota` (default)
- `vendor`

`dakota` denotes Dakota's internal finite differencing algorithm and `vendor` denotes the finite differencing algorithm supplied by the iterator package in use (DOT, CONMIN, NPSOL, NL2SOL, NLSSOL, and OPT++ each have their own internal finite differencing routines). The `dakota` routine is the default since it can execute in parallel and exploit the concurrency in finite difference evaluations (see Exploiting Parallelism in the Users Manual [[4] "Adams et al., 2010"]).

However, the `vendor` setting can be desirable in some cases since certain libraries will modify their algorithm when the finite differencing is performed internally. Since the selection of the `dakota` routine hides the use of finite differencing from the optimizers (the optimizers are configured to accept user-supplied gradients, which some algorithms assume to be of analytic accuracy), the potential exists for the `vendor` setting to trigger the

use of an algorithm more optimized for the higher expense and/or lower accuracy of finite-differencing. For example, NPSOL uses gradients in its line search when in user-supplied gradient mode (since it assumes they are inexpensive), but uses a value-based line search procedure when internally finite differencing. The use of a value-based line search will often reduce total expense in serial operations. However, in parallel operations, the use of gradients in the NPSOL line search (user-supplied gradient mode) provides excellent load balancing without need to resort to speculative optimization approaches.

In summary, then, the `dakota` routine is preferred for parallel optimization, and the `vendor` routine may be preferred for serial optimization in special cases.

dakota

- Keywords Area
- responses
- mixed_gradients
- dakota

(Default) Use internal Dakota finite differences algorithm

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		ignore_bounds	Do not respect bounds when computing gradients or Hessians
	Optional*(Choose One)*	**Group 1**	relative	Scale step size by the parameter value
			absolute	Do not scale step-size
			bounds	Scale step-size by the domain of the parameter

Description

The `dakota` routine is the default since it can execute in parallel and exploit the concurrency in finite difference evaluations (see Exploiting Parallelism in the Users Manual [[4] "Adams et al., 2010"]).

When the `method_source` is `dakota`, the user may also specify the type of scaling desired when determining the finite difference step size. The choices are `absolute`, `bounds`, and `relative`. For `absolute`, the step size will be applied as is. For `bounds`, it will be scaled by the range of each parameter. For `relative`, it will be scaled by the parameter value.

ignore_bounds

- Keywords Area
- responses
- mixed_gradients
- dakota
- ignore_bounds

Do not respect bounds when computing gradients or Hessians

Specification

Alias: none
Argument(s): none

Description

When Dakota computes gradients or Hessians by finite differences and the variables in question have bounds, it by default chooses finite-differencing steps that keep the variables within their specified bounds. Older versions of Dakota generally ignored bounds when computing finite differences. To restore the older behavior, one can add keyword `ignore_bounds` to the `response` specification when `method_source dakota` (or just `dakota`) is also specified.

In forward difference or backward difference computations, honoring bounds is straightforward.

To honor bounds when approximating $\partial f/\partial x_i$, i.e., component i of the gradient of f, by central differences, Dakota chooses two steps h_1 and h_2 with $h_1 \neq h_2$, such that $x + h_1 e_i$ and $x + h_2 e_i$ both satisfy the bounds, and then computes

$$\frac{\partial f}{\partial x_i} \cong \frac{h_2^2(f_1 - f_0) - h_1^2(f_2 - f_0)}{h_1 h_2 (h_2 - h_1)},$$

with $f_0 = f(x)$, $f_1 = f(x + h_1 e_i)$, and $f_2 = f(x + h_2 e_i)$.

relative

- Keywords Area
- responses
- mixed_gradients
- dakota
- relative

Scale step size by the parameter value

Specification

Alias: none
Argument(s): none

Description

Scale step size by the parameter value

absolute

- Keywords Area
- responses
- mixed_gradients
- dakota
- absolute

Do not scale step-size

Specification

Alias: none
Argument(s): none

Description

Do not scale step-size

bounds

- Keywords Area
- responses
- mixed_gradients
- dakota
- bounds

Scale step-size by the domain of the parameter

Specification

Alias: none
Argument(s): none

Description

Scale step-size by the domain of the parameter

vendor

- Keywords Area
- responses
- mixed_gradients
- vendor

Use non-Dakota fd algorithm

Specification

Alias: none
Argument(s): none

Description

See parent page for usage notes.

interval_type

- Keywords Area
- responses
- mixed_gradients
- interval_type

Specify how to compute gradients and hessians

Specification

Alias: none
Argument(s): none

Description

The `interval_type` setting is used to select between `forward` and `central` differences in the numerical gradient calculations. The `dakota`, DOT `vendor`, and OPT++ `vendor` routines have both forward and central differences available, the CONMIN and NL2SOL `vendor` routines support forward differences only, and the NPSOL and NLSSOL `vendor` routines start with forward differences and automatically switch to central differences as the iteration progresses (the user has no control over this). The following forward difference expression

$$\nabla f(\mathbf{x}) \cong \frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h}$$

and the following central difference expression

$$\nabla f(\mathbf{x}) \cong \frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x} - h\mathbf{e}_i)}{2h}$$

are used to estimate the ith component of the gradient vector.

forward

- Keywords Area
- responses
- mixed_gradients
- forward

Use forward differences

Specification

Alias: none
Argument(s): none

Description

See parent page for usage notes.

central

- Keywords Area
- responses
- mixed_gradients
- central

Use central differences

Specification

Alias: none
Argument(s): none

Description

See parent page for usage notes.

fd_step_size

- Keywords Area
- responses
- mixed_gradients
- fd_step_size

Step size used when computing gradients and Hessians

Specification

Alias: fd_gradient_step_size
Argument(s): REALLIST

Description

`fd_gradient_step_size` specifies the relative finite difference step size to be used in the computations. Either a single value may be entered for use with all parameters, or a list of step sizes may be entered, one for each parameter.

The latter option of a list of step sizes is only valid for use with the Dakota finite differencing routine. For Dakota with an interval scaling type of `absolute`, the differencing interval will be `fd_gradient_step_size`.

For Dakota with and interval scaling type of `bounds`, the differencing intervals are computed by multiplying `fd_gradient_step_size` with the range of the parameter. For Dakota (with an interval scaling type of `relative`), DOT, CONMIN, and OPT++, the differencing intervals are computed by multiplying the `fd_gradient_step_size` with the current parameter value. In this case, a minimum absolute differencing interval is needed when the current parameter value is close to zero. This prevents finite difference intervals for the parameter which are too small to distinguish differences in the response quantities being computed. Dakota, DOT, CONMIN, and OPT++ all use `.01*fd_gradient_step_size` as their minimum absolute differencing interval. With a `fd_gradient_step_size = .001`, for example, Dakota, DOT, CONMIN, and OPT++ will use intervals of .001∗current value with a minimum interval of 1.e-5. NPSOL and NLSSOL use a different formula for their finite difference intervals: `fd_gradient_step_size*(1+|current parameter value|)`. This definition has the advantage of eliminating the need for a minimum absolute differencing interval since the interval no longer goes to zero as the current parameter value goes to zero.

5.6.9 numerical_gradients

- Keywords Area
- responses
- numerical_gradients

Gradients are needed and will be approximated by finite differences

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		method_source	Specify which finite difference routine is used

	Optional*(Choose One)*	**Group 1**	dakota	(Default) Use internal Dakota finite differences algorithm
			vendor	Use non-Dakota fd algorithm
	Optional		interval_type	Specify how to compute gradients and hessians
	Optional*(Choose One)*	**Group 2**	forward	Use forward differences
			central	Use central differences
	Optional		fd_step_size	Step size used when computing gradients and Hessians

Description

The `numerical_gradients` specification means that gradient information is needed and will be computed with finite differences using either the native or one of the vendor finite differencing routines.

See Also

These keywords may also be of interest:

- no_gradients
- analytic_gradients
- mixed_gradients

method_source

- Keywords Area
- responses
- numerical_gradients
- method_source

Specify which finite difference routine is used

Specification

Alias: none
Argument(s): none

Description

The `method_source` setting specifies the source of the finite differencing routine that will be used to compute the numerical gradients:

- `dakota` (default)
- `vendor`

`dakota` denotes Dakota's internal finite differencing algorithm and `vendor` denotes the finite differencing algorithm supplied by the iterator package in use (DOT, CONMIN, NPSOL, NL2SOL, NLSSOL, and OPT++ each have their own internal finite differencing routines). The `dakota` routine is the default since it can execute in parallel and exploit the concurrency in finite difference evaluations (see Exploiting Parallelism in the Users Manual [[4] "Adams et al., 2010"]).

However, the `vendor` setting can be desirable in some cases since certain libraries will modify their algorithm when the finite differencing is performed internally. Since the selection of the `dakota` routine hides the use of finite differencing from the optimizers (the optimizers are configured to accept user-supplied gradients, which some algorithms assume to be of analytic accuracy), the potential exists for the `vendor` setting to trigger the use of an algorithm more optimized for the higher expense and/or lower accuracy of finite-differencing. For example, NPSOL uses gradients in its line search when in user-supplied gradient mode (since it assumes they are inexpensive), but uses a value-based line search procedure when internally finite differencing. The use of a value-based line search will often reduce total expense in serial operations. However, in parallel operations, the use of gradients in the NPSOL line search (user-supplied gradient mode) provides excellent load balancing without need to resort to speculative optimization approaches.

In summary, then, the `dakota` routine is preferred for parallel optimization, and the `vendor` routine may be preferred for serial optimization in special cases.

dakota

- Keywords Area
- responses
- numerical_gradients
- dakota

(Default) Use internal Dakota finite differences algorithm

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		ignore_bounds	Do not respect bounds when computing gradients or Hessians

	Optional*(Choose One)*	**Group 1**	relative	Scale step size by the parameter value
			absolute	Do not scale step-size
			bounds	Scale step-size by the domain of the parameter

Description

The `dakota` routine is the default since it can execute in parallel and exploit the concurrency in finite difference evaluations (see Exploiting Parallelism in the Users Manual [[4] "Adams et al., 2010"]).

When the `method_source` is `dakota`, the user may also specify the type of scaling desired when determining the finite difference step size. The choices are `absolute`, `bounds`, and `relative`. For `absolute`, the step size will be applied as is. For `bounds`, it will be scaled by the range of each parameter. For `relative`, it will be scaled by the parameter value.

ignore_bounds

- Keywords Area
- responses
- numerical_gradients
- dakota
- ignore_bounds

Do not respect bounds when computing gradients or Hessians

Specification

Alias: none

Argument(s): none

Description

When Dakota computes gradients or Hessians by finite differences and the variables in question have bounds, it by default chooses finite-differencing steps that keep the variables within their specified bounds. Older versions of Dakota generally ignored bounds when computing finite differences. To restore the older behavior, one can add keyword `ignore_bounds` to the `response` specification when `method_source dakota` (or just `dakota`) is also specified.

In forward difference or backward difference computations, honoring bounds is straightforward.

To honor bounds when approximating $\partial f/\partial x_i$, i.e., component i of the gradient of f, by central differences, Dakota chooses two steps h_1 and h_2 with $h_1 \neq h_2$, such that $x + h_1 e_i$ and $x + h_2 e_i$ both satisfy the bounds, and then computes

$$\frac{\partial f}{\partial x_i} \cong \frac{h_2^2(f_1 - f_0) - h_1^2(f_2 - f_0)}{h_1 h_2 (h_2 - h_1)},$$

with $f_0 = f(x)$, $f_1 = f(x + h_1 e_i)$, and $f_2 = f(x + h_2 e_i)$.

relative

- Keywords Area
- responses
- numerical_gradients
- dakota
- relative

Scale step size by the parameter value

Specification

Alias: none
Argument(s): none

Description

Scale step size by the parameter value

absolute

- Keywords Area
- responses
- numerical_gradients
- dakota
- absolute

Do not scale step-size

Specification

Alias: none
Argument(s): none

Description

Do not scale step-size

bounds

- Keywords Area
- responses
- numerical_gradients
- dakota
- bounds

Scale step-size by the domain of the parameter

Specification

Alias: none
Argument(s): none

Description

Scale step-size by the domain of the parameter

vendor

- Keywords Area
- responses
- numerical_gradients
- vendor

Use non-Dakota fd algorithm

Specification

Alias: none
Argument(s): none

Description

See parent page for usage notes.

interval_type

- Keywords Area
- responses
- numerical_gradients
- interval_type

Specify how to compute gradients and hessians

Specification

Alias: none
Argument(s): none

Description

The `interval_type` setting is used to select between `forward` and `central` differences in the numerical gradient calculations. The `dakota`, DOT `vendor`, and OPT++ `vendor` routines have both forward and central differences available, the CONMIN and NL2SOL `vendor` routines support forward differences only, and the NPSOL and NLSSOL `vendor` routines start with forward differences and automatically switch to central differences as the iteration progresses (the user has no control over this). The following forward difference expression

$$\nabla f(\mathbf{x}) \cong \frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h}$$

and the following central difference expression

$$\nabla f(\mathbf{x}) \cong \frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x} - h\mathbf{e}_i)}{2h}$$

are used to estimate the ith component of the gradient vector.

forward

- Keywords Area
- responses
- numerical_gradients
- forward

Use forward differences

Specification

Alias: none
Argument(s): none

Description

See parent page for usage notes.

central

- Keywords Area
- responses
- numerical_gradients
- central

Use central differences

Specification

Alias: none
Argument(s): none

Description

See parent page for usage notes.

fd_step_size

- Keywords Area
- responses
- numerical_gradients
- fd_step_size

Step size used when computing gradients and Hessians

Specification

Alias: fd_gradient_step_size
Argument(s): REALLIST

Description

`fd_gradient_step_size` specifies the relative finite difference step size to be used in the computations. Either a single value may be entered for use with all parameters, or a list of step sizes may be entered, one for each parameter.

The latter option of a list of step sizes is only valid for use with the Dakota finite differencing routine. For Dakota with an interval scaling type of `absolute`, the differencing interval will be `fd_gradient_step_size`.

For Dakota with and interval scaling type of `bounds`, the differencing intervals are computed by multiplying `fd_gradient_step_size` with the range of the parameter. For Dakota (with an interval scaling type of `relative`), DOT, CONMIN, and OPT++, the differencing intervals are computed by multiplying the `fd_gradient_step_size` with the current parameter value. In this case, a minimum absolute differencing interval is needed when the current parameter value is close to zero. This prevents finite difference intervals for the parameter which are too small to distinguish differences in the response quantities being computed. Dakota, DOT, CONMIN, and OPT++ all use `.01*fd_gradient_step_size` as their minimum absolute differencing interval. With a `fd_gradient_step_size = .001`, for example, Dakota, DOT, CONMIN, and OPT++ will use intervals of .001∗current value with a minimum interval of 1.e-5. NPSOL and NLSSOL use a different formula for their finite difference intervals: `fd_gradient_step_size*(1+|current parameter value|)`. This definition has the advantage of eliminating the need for a minimum absolute differencing interval since the interval no longer goes to zero as the current parameter value goes to zero.

5.6.10 no_hessians

- Keywords Area
- responses
- no_hessians

Hessians will not be used

Specification

Alias: none
Argument(s): none

Description

The `no_hessians` specification means that the method does not require Dakota to manage the computation of any Hessian information. Therefore, it will neither be retrieved from the simulation nor computed by Dakota. The `no_hessians` keyword is a complete specification for this case. Note that, in some cases, Hessian information may still be being approximated internal to an algorithm (e.g., within a quasi-Newton optimizer such as `optpp_q_newton`); however, Dakota has no direct involvement in this process and the responses specification need not include it.

See Also

These keywords may also be of interest:

- numerical_hessians
- quasi_hessians
- analytic_hessians
- mixed_hessians

5.6.11 numerical_hessians

- Keywords Area
- responses
- numerical_hessians

Hessians are needed and will be approximated by finite differences

Specification

Alias: none
Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		fd_step_size	Step size used when computing gradients and Hessians

	Optional*(Choose One)*	**Group 1**	relative	Scale step size by the parameter value
			absolute	Do not scale step-size
			bounds	Scale step-size by the domain of the parameter
	Optional*(Choose One)*	**Group 2**	forward	Use forward differences
			central	Use central differences

Description

The `numerical_hessians` specification means that Hessian information is needed and will be computed with finite differences using either first-order gradient differencing (for the cases of `analytic_gradients` or for the functions identified by `id_analytic_gradients` in the case of `mixed_gradients`) or first- or second-order function value differencing (all other gradient specifications). In the former case, the following expression

$$\nabla^2 f(\mathbf{x})_i \cong \frac{\nabla f(\mathbf{x} + h\mathbf{e}_i) - \nabla f(\mathbf{x})}{h}$$

estimates the ith Hessian column, and in the latter case, the following expressions

$$\nabla^2 f(\mathbf{x})_{i,j} \cong \frac{f(\mathbf{x} + h_i\mathbf{e}_i + h_j\mathbf{e}_j) - f(\mathbf{x} + h_i\mathbf{e}_i) - f(\mathbf{x} - h_j\mathbf{e}_j) + f(\mathbf{x})}{h_i h_j}$$

and

$$\nabla^2 f(\mathbf{x})_{i,j} \cong \frac{f(\mathbf{x} + h\mathbf{e}_i + h\mathbf{e}_j) - f(\mathbf{x} + h\mathbf{e}_i - h\mathbf{e}_j) - f(\mathbf{x} - h\mathbf{e}_i + h\mathbf{e}_j) + f(\mathbf{x} - h\mathbf{e}_i - h\mathbf{e}_j)}{4h^2}$$

provide first- and second-order estimates of the ij^{th} Hessian term. Prior to Dakota 5.0, Dakota always used second-order estimates. In Dakota 5.0 and newer, the default is to use first-order estimates (which honor bounds on the variables and require only about a quarter as many function evaluations as do the second-order estimates), but specifying `central` after `numerical_hessians` causes Dakota to use the old second-order estimates, which do not honor bounds. In optimization algorithms that use Hessians, there is little reason to use second-order differences in computing Hessian approximations.

See Also

These keywords may also be of interest:

- no_hessians
- quasi_hessians
- analytic_hessians
- mixed_hessians

fd_step_size

- Keywords Area
- responses
- numerical_hessians
- fd_step_size

Step size used when computing gradients and Hessians

Specification

Alias: fd_hessian_step_size
Argument(s): REALLIST

Description

`fd_gradient_step_size` specifies the relative finite difference step size to be used in the computations. Either a single value may be entered for use with all parameters, or a list of step sizes may be entered, one for each parameter.

The latter option of a list of step sizes is only valid for use with the Dakota finite differencing routine. For Dakota with an interval scaling type of `absolute`, the differencing interval will be `fd_gradient_step_size`.

For Dakota with and interval scaling type of `bounds`, the differencing intervals are computed by multiplying `fd_gradient_step_size` with the range of the parameter. For Dakota (with an interval scaling type of `relative`), DOT, CONMIN, and OPT++, the differencing intervals are computed by multiplying the `fd_gradient_step_size` with the current parameter value. In this case, a minimum absolute differencing interval is needed when the current parameter value is close to zero. This prevents finite difference intervals for the parameter which are too small to distinguish differences in the response quantities being computed. Dakota, DOT, CONMIN, and OPT++ all use `.01*fd_gradient_step_size` as their minimum absolute differencing interval. With a `fd_gradient_step_size = .001`, for example, Dakota, DOT, CONMIN, and OPT++ will use intervals of .001∗current value with a minimum interval of 1.e-5. NPSOL and NLSSOL use a different formula for their finite difference intervals: `fd_gradient_step_size*(1+|current parameter value|)`. This definition has the advantage of eliminating the need for a minimum absolute differencing interval since the interval no longer goes to zero as the current parameter value goes to zero.

relative

- Keywords Area
- responses
- numerical_hessians
- relative

Scale step size by the parameter value

Specification

Alias: none
Argument(s): none

Description

Scale step size by the parameter value

absolute

- Keywords Area
- responses
- numerical_hessians
- absolute

Do not scale step-size

Specification

Alias: none
Argument(s): none

Description

Do not scale step-size

bounds

- Keywords Area
- responses
- numerical_hessians
- bounds

Scale step-size by the domain of the parameter

Specification

Alias: none
Argument(s): none

Description

Scale step-size by the domain of the parameter

forward

- Keywords Area
- responses
- numerical_hessians
- forward

Use forward differences

Specification

Alias: none
Argument(s): none

Description

See parent page for usage notes.

central

- Keywords Area
- responses
- numerical_hessians
- central

Use central differences

Specification

Alias: none
Argument(s): none

Description

See parent page for usage notes.

5.6.12 quasi_hessians

- Keywords Area
- responses
- quasi_hessians

Hessians are needed and will be approximated by secant updates (BFGS or SR1) from a series of gradient evaluations

Specification

Alias: none
Argument(s): none

	Required/-Optional	**Description of Group**	**Dakota Keyword**	**Dakota Keyword Description**
	Required*(Choose One)*	**Group 1**	bfgs	Use BFGS method to compute quasi-hessians

			sr1	Use the Symmetric Rank 1 update method to compute quasi-Hessians

Description

The `quasi_hessians` specification means that Hessian information is needed and will be approximated using secant updates (sometimes called "quasi-Newton updates", though any algorithm that approximates Newton's method is a quasi-Newton method).

Compared to finite difference numerical Hessians, secant approximations do not expend additional function evaluations in estimating all of the second-order information for every point of interest. Rather, they accumulate approximate curvature information over time using the existing gradient evaluations.

The supported secant approximations include the Broyden-Fletcher-Goldfarb-Shanno (BFGS) update (specified with the keyword `bfgs`) and the Symmetric Rank 1 (SR1) update (specified with the keyword `sr1`).

See Also

These keywords may also be of interest:

- no_hessians
- numerical_hessians
- analytic_hessians
- mixed_hessians

bfgs

- Keywords Area
- responses
- quasi_hessians
- bfgs

Use BFGS method to compute quasi-hessians

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		damped	Numerical safeguarding for BFGS updates

Description

Broyden-Fletcher-Goldfarb-Shanno (BFGS) update will be used to compute quasi-Hessians.

$$B_{k+1} = B_k - \frac{B_k s_k s_k^T B_k}{s_k^T B_k s_k} + \frac{y_k y_k^T}{y_k^T s_k}$$

where B_k is the kth approximation to the Hessian, $s_k = x_{k+1} - x_k$ is the step and $y_k = \nabla f_{k+1} - \nabla f_k$ is the corresponding yield in the gradients.

Notes

- Initial scaling of $\frac{y_k^T y_k}{y_k^T s_k} I$ is used for B_0 prior to the first update.
- Numerical safeguarding is used to protect against numerically small denominators within the updates.
- This safeguarding skips the update if $|y_k^T s_k| < 10^{-6} s_k^T B_k s_k$
- Additional safeguarding can be added using the `damped` option, which utilizes an alternative damped BFGS update when the curvature condition $y_k^T s_k > 0$ is nearly violated.

damped

- Keywords Area
- responses
- quasi_hessians
- bfgs
- damped

Numerical safeguarding for BFGS updates

Specification

Alias: none

Argument(s): none

Description

See parent page.

sr1

- Keywords Area
- responses
- quasi_hessians
- sr1

Use the Symmetric Rank 1 update method to compute quasi-Hessians

Specification

Alias: none
Argument(s): none

Description

The Symmetric Rank 1 (SR1) update (specified with the keyword `sr1`) will be used to compute quasi-Hessians.

$$B_{k+1} = B_k + \frac{(y_k - B_k s_k)(y_k - B_k s_k)^T}{(y_k - B_k s_k)^T s_k}$$

where B_k is the kth approximation to the Hessian, $s_k = x_{k+1} - x_k$ is the step and $y_k = \nabla f_{k+1} - \nabla f_k$ is the corresponding yield in the gradients.

Notes

- Initial scaling of $\frac{y_k^T y_k}{y_k^T s_k} I$ is used for B_0 prior to the first update.
- Numerical safeguarding is used to protect against numerically small denominators within the updates.
- This safeguarding skips the update if $|(y_k - B_k s_k)^T s_k| < 10^{-6} ||s_k||_2 ||y_k - B_k s_k||_2$

5.6.13 analytic_hessians

- Keywords Area
- responses
- analytic_hessians

Hessians are needed and are available directly from the analysis driver

Specification

Alias: none
Argument(s): none

Description

The `analytic_hessians` specification means that Hessian information is available directly from the simulation. The simulation must return the Hessian data in the Dakota format (enclosed in double brackets; see Dakota File Data Formats in Users Manual [[4] "Adams et al., 2010"]) for the case of file transfer of data. The `analytic_hessians` keyword is a complete specification for this case.

See Also

These keywords may also be of interest:

- no_hessians
- numerical_hessians
- quasi_hessians
- mixed_hessians

5.6.14 mixed_hessians

- Keywords Area
- responses
- mixed_hessians

Hessians are needed and will be obtained from a mix of numerical, analytic, and "quasi" sources

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		id_numerical_-hessians	Identify which numerical-Hessian corresponds to which response
	Optional*(Choose One)*	**Group 1**	relative	Scale step size by the parameter value
			absolute	Do not scale step-size
			bounds	Scale step-size by the domain of the parameter
	Optional*(Choose One)*	**Group 2**	forward	Use forward differences
			central	Use central differences
	Optional		id_quasi_hessians	Identify which quasi-Hessian corresponds to which response
	Optional		id_analytic_-hessians	Identify which analytical Hessian corresponds to which response

Description

Hessian availability must be specified with either `no_hessians`, `numerical_hessians`, `quasi_hessians`, `analytic_hessians`, or `mixed_hessians`.

The `mixed_hessians` specification means that some Hessian information is available directly from the simulation (analytic) whereas the rest will have to be estimated by finite differences (numerical) or approximated by secant updating. As for mixed gradients, this specification allows the user to make use of as much analytic information as is available and then estimate/approximate the rest.

The `id_analytic_hessians` list specifies by number the functions which have analytic Hessians, and the `id_numerical_hessians` and `id_quasi_hessians` lists specify by number the functions which must

use numerical Hessians and secant Hessian updates, respectively. Each function identifier, from 1 through the total number of functions, must appear once and only once within the union of the `id_analytic_hessians`, `id_numerical_hessians`, and `id_quasi_hessians` lists.

The `fd_hessian_step_size` and `bfgs`, `damped bfgs`, or `sr1` secant update selections are as described previously in responses and pertain to those functions listed by the `id_numerical_hessians` and `id_quasi_hessians` lists.

See Also

These keywords may also be of interest:

- no_hessians
- numerical_hessians
- quasi_hessians
- analytic_hessians

id_numerical_hessians

- Keywords Area
- responses
- mixed_hessians
- id_numerical_hessians

Identify which numerical-Hessian corresponds to which response

Topics

This keyword is related to the topics:

- objective_function_pointer

Specification

Alias: none

Argument(s): INTEGERLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		fd_step_size	Step size used when computing gradients and Hessians

Description

The `id_analytic_hessians` list specifies by number the functions which have analytic Hessians, and the `id_numerical_hessians` and `id_quasi_hessians` lists specify by number the functions which must use numerical Hessians and secant Hessian updates, respectively. Each function identifier, from 1 through the total number of functions, must appear once and only once within the union of the `id_analytic_hessians`, `id_numerical_hessians`, and `id_quasi_hessians` lists.

See Also

These keywords may also be of interest:

- id_analytic_hessians
- id_quasi_hessians

fd_step_size

- Keywords Area
- responses
- mixed_hessians
- id_numerical_hessians
- fd_step_size

Step size used when computing gradients and Hessians

Specification

Alias: fd_hessian_step_size
Argument(s): REALLIST

Description

`fd_gradient_step_size` specifies the relative finite difference step size to be used in the computations. Either a single value may be entered for use with all parameters, or a list of step sizes may be entered, one for each parameter.

The latter option of a list of step sizes is only valid for use with the Dakota finite differencing routine. For Dakota with an interval scaling type of `absolute`, the differencing interval will be `fd_gradient_step_size`.

For Dakota with and interval scaling type of `bounds`, the differencing intervals are computed by multiplying `fd_gradient_step_size` with the range of the parameter. For Dakota (with an interval scaling type of `relative`), DOT, CONMIN, and OPT++, the differencing intervals are computed by multiplying the `fd_gradient_step_size` with the current parameter value. In this case, a minimum absolute differencing interval is needed when the current parameter value is close to zero. This prevents finite difference intervals for the parameter which are too small to distinguish differences in the response quantities being computed. Dakota, DOT, CONMIN, and OPT++ all use `.01*fd_gradient_step_size` as their minimum absolute differencing interval. With a `fd_gradient_step_size = .001`, for example, Dakota, DOT, CONMIN, and OPT++ will use intervals of .001∗current value with a minimum interval of 1.e-5. NPSOL and NLSSOL use a different formula for their finite difference intervals: `fd_gradient_step_size*(1+|current parameter value|)`. This definition has the advantage of eliminating the need for a minimum absolute differencing interval since the interval no longer goes to zero as the current parameter value goes to zero.

relative

- Keywords Area
- responses
- mixed_hessians
- relative

Scale step size by the parameter value

Specification

Alias: none
Argument(s): none

Description

Scale step size by the parameter value

absolute

- Keywords Area
- responses
- mixed_hessians
- absolute

Do not scale step-size

Specification

Alias: none
Argument(s): none

Description

Do not scale step-size

bounds

- Keywords Area
- responses
- mixed_hessians
- bounds

Scale step-size by the domain of the parameter

Specification

Alias: none
Argument(s): none

Description

Scale step-size by the domain of the parameter

forward

- Keywords Area
- responses
- mixed_hessians
- forward

Use forward differences

Specification

Alias: none
Argument(s): none

Description

See parent page for usage notes.

central

- Keywords Area
- responses
- mixed_hessians
- central

Use central differences

Specification

Alias: none
Argument(s): none

Description

See parent page for usage notes.

id_quasi_hessians

- Keywords Area
- responses
- mixed_hessians
- id_quasi_hessians

Identify which quasi-Hessian corresponds to which response

Topics

This keyword is related to the topics:

- objective_function_pointer

Specification

Alias: none

Argument(s): INTEGERLIST

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Required*(Choose One)*	**Group 1**	bfgs	Use BFGS method to compute quasi-hessians
			sr1	Use the Symmetric Rank 1 update method to compute quasi-Hessians

Description

The `id_analytic_hessians` list specifies by number the functions which have analytic Hessians, and the `id_numerical_hessians` and `id_quasi_hessians` lists specify by number the functions which must use numerical Hessians and secant Hessian updates, respectively. Each function identifier, from 1 through the total number of functions, must appear once and only once within the union of the `id_analytic_hessians`, `id_numerical_hessians`, and `id_quasi_hessians` lists.

See Also

These keywords may also be of interest:

- id_numerical_hessians
- id_analytic_hessians

bfgs

- Keywords Area
- responses
- mixed_hessians
- id_quasi_hessians
- bfgs

Use BFGS method to compute quasi-hessians

Specification

Alias: none

Argument(s): none

	Required/-Optional	Description of Group	Dakota Keyword	Dakota Keyword Description
	Optional		damped	Numerical safeguarding for BFGS updates

Description

Broyden-Fletcher-Goldfarb-Shanno (BFGS) update will be used to compute quasi-Hessians.

$$B_{k+1} = B_k - \frac{B_k s_k s_k^T B_k}{s_k^T B_k s_k} + \frac{y_k y_k^T}{y_k^T s_k}$$

where B_k is the kth approximation to the Hessian, $s_k = x_{k+1} - x_k$ is the step and $y_k = \nabla f_{k+1} - \nabla f_k$ is the corresponding yield in the gradients.

Notes

- Initial scaling of $\frac{y_k^T y_k}{y_k^T s_k} I$ is used for B_0 prior to the first update.
- Numerical safeguarding is used to protect against numerically small denominators within the updates.
- This safeguarding skips the update if $|y_k^T s_k| < 10^{-6} s_k^T B_k s_k$
- Additional safeguarding can be added using the `damped` option, which utilizes an alternative damped BFGS update when the curvature condition $y_k^T s_k > 0$ is nearly violated.

damped

- Keywords Area
- responses
- mixed_hessians
- id_quasi_hessians
- bfgs
- damped

Numerical safeguarding for BFGS updates

Specification

Alias: none
Argument(s): none

Description

See parent page.

sr1

- Keywords Area
- responses
- mixed_hessians
- id_quasi_hessians
- sr1

Use the Symmetric Rank 1 update method to compute quasi-Hessians

Specification

Alias: none
Argument(s): none

Description

The Symmetric Rank 1 (SR1) update (specified with the keyword `sr1`) will be used to compute quasi-Hessians.

$$B_{k+1} = B_k + \frac{(y_k - B_k s_k)(y_k - B_k s_k)^T}{(y_k - B_k s_k)^T s_k}$$

where B_k is the kth approximation to the Hessian, $s_k = x_{k+1} - x_k$ is the step and $y_k = \nabla f_{k+1} - \nabla f_k$ is the corresponding yield in the gradients.

Notes

- Initial scaling of $\frac{y_k^T y_k}{y_k^T s_k} I$ is used for B_0 prior to the first update.
- Numerical safeguarding is used to protect against numerically small denominators within the updates.
- This safeguarding skips the update if $|(y_k - B_k s_k)^T s_k| < 10^{-6} ||s_k||_2 ||y_k - B_k s_k||_2$

id_analytic_hessians

- Keywords Area
- responses
- mixed_hessians
- id_analytic_hessians

Identify which analytical Hessian corresponds to which response

Topics

This keyword is related to the topics:

- objective_function_pointer

Specification

Alias: none
Argument(s): INTEGERLIST

Description

The `id_analytic_hessians` list specifies by number the functions which have analytic Hessians, and the `id_numerical_hessians` and `id_quasi_hessians` lists specify by number the functions which must use numerical Hessians and secant Hessian updates, respectively. Each function identifier, from 1 through the total number of functions, must appear once and only once within the union of the `id_analytic_hessians`, `id_numerical_hessians`, and `id_quasi_hessians` lists.

See Also

These keywords may also be of interest:

- id_numerical_hessians
- id_quasi_hessians

Bibliography

[1] Computational investigations of low-discrepancy sequences. *ACM Transactions on Mathematical Software*, 23(2):266–294, 1997. 1038, 1111, 1112

[2] M.A. Abramson, C. Audet, G. Couture, J.E. Dennis, Jr., S. Le Digabel, and C. Tribes. The NOMAD project. Software available at `http://www.gerad.ca/nomad`. 470

[3] B. M. Adams, L. E. Bauman, W. J. Bohnhoff, K. R. Dalbey, J. P. Eddy, M. S. Ebeida, M. S. Eldred, P. D. Hough, K. T. Hu, J. D. Jakeman, L. P. Swiler, Stephens J. A., D. M. Vigil, and T. M. Wildey. Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 6.0 developers manual. Technical Report SAND2014-5014, Sandia National Laboratories, Albuquerque, NM, Updated May 2014. Available online from `http://dakota.sandia.gov/documentation.html`. 7

[4] B. M. Adams, L. E. Bauman, W. J. Bohnhoff, K. R. Dalbey, J. P. Eddy, M. S. Ebeida, M. S. Eldred, P. D. Hough, K. T. Hu, J. D. Jakeman, L. P. Swiler, Stephens J. A., D. M. Vigil, and T. M. Wildey. Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 6.0 users manual. Technical Report SAND2014-4633, Sandia National Laboratories, Albuquerque, NM, Updated May 2014. Available online from `http://dakota.sandia.gov/documentation.html`. 7, 13, 14, 23, 106, 172, 175, 176, 180, 183, 188, 191, 192, 202, 203, 264, 317, 667, 752, 1061, 1063, 1221, 1224, 1339, 1348, 1349, 1357, 1360, 1369, 1370, 1371, 1373, 1375, 1384, 1388, 1419, 1422, 1423, 1430, 1431, 1443

[5] B. M. Adams, L. E. Bauman, W. J. Bohnhoff, K. R. Dalbey, J. P. Eddy, M. S. Ebeida, M. S. Eldred, P. D. Hough, K. T. Hu, J. D. Jakeman, L. P. Swiler, Stephens J. A., D. M. Vigil, and T. M. Wildey. Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 6.0 reference manual. Technical Report SAND2014-5015, Sandia National Laboratories, Albuquerque, NM, Updated May 2014. Available online from `http://dakota.sandia.gov/documentation.html`. 107, 110

[6] G. Anderson and P. Anderson. *The UNIX C Shell Field Guide*. Prentice-Hall, Englewood Cliffs, NJ, 1986. 10

[7] J. S. Arora. *Introduction to Optimum Design*. McGraw-Hill, New York, 1989. 115

[8] C. Audet, S. Le Digabel, and C. Tribes. NOMAD user guide. Technical Report G-2009-37, Les cahiers du GERAD, 2009. 470

[9] J.-P. Berrut and L. N. Trefethen. Barycentric lagrange interpolation. *SIAM Review*, 46(3):501–517, 2004. 112

[10] B. J. Bichon, M. S. Eldred, L. P. Swiler, S. Mahadevan, and J. M. McFarland. Multimodal reliability assessment for complex engineering applications using efficient global optimization. In *Proceedings of the*

48th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference (9th AIAA Non-Deterministic Approaches Conference), number AIAA-2007-1946, Honolulu, HI, April 2007. 1088

[11] B. J. Bichon, M. S. Eldred, L. P. Swiler, S. Mahadevan, and J. M. McFarland. Efficient global reliability analysis for nonlinear implicit performance functions. *AIAA Journal*, 46(10):2459–2468, 2008. 1088

[12] K. Breitung. Asymptotic approximation for multinormal integrals. *J. Eng. Mech., ASCE*, 110(3):357–366, 1984. 1063

[13] R. H. Byrd, R. B. Schnabel, and G. A. Schultz. Parallel quasi-newton methods for unconstrained optimization. *Mathematical Programming*, 42:273–306, 1988. 157

[14] K. J. Chang, R. T. Haftka, G. L. Giles, and P.-J. Kao. Sensitivity-based scaling for approximating structural response. *J. Aircraft*, 30:283–288, 1993. 1192, 1210

[15] A. R. Conn, N. I. M. Gould, and P. L. Toint. *Trust-Region Methods*. MPS-SIAM Series on Optimization, SIAM-MPS, Philadelphia, 2000. 206, 207, 208

[16] A. Der Kiureghian and P. L. Liu. Structural reliability under incomplete information. *J. Eng. Mech., ASCE*, 112(EM-1):85–104, 1986. 111

[17] Q. Du, V. Faber, and M. Gunzburger. Centroidal voronoi tessellations: Applications and algorithms. *SIAM Review*, 41:637–676, 1999. 1031

[18] J. E. Eddy and K. Lewis. Effective generation of pareto sets using genetic programming. In *Proceedings of ASME Design Engineering Technical Conference*, 2001. 133

[19] A. S. El-Bakry, R. A. Tapia, T. Tsuchiya, and Y. Zhang. On the formulation and theory of the newton interior-point method for nonlinear programming. *Journal of Optimization Theory and Applications*, 89:507–541, 1996. 377, 394, 411, 428

[20] M. S. Eldred, H. Agarwal, V. M. Perez, S. F. Wojtkiewicz, Jr., and J. E. Renaud. Investigation of reliability method formulations in DAKOTA/UQ. *Structure & Infrastructure Engineering: Maintenance, Management, Life-Cycle Design & Performance*, 3(3):199–213, 2007. 108

[21] M. S. Eldred and B. J. Bichon. Second-order reliability formulations in DAKOTA/UQ. In *Proceedings of the 47th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, number AIAA-2006-1828, Newport, RI, May 1–4 2006. 1063

[22] M. S. Eldred and D. M. Dunlavy. Formulations for surrogate-based optimization with data fit, multifidelity, and reduced-order models. In *Proceedings of the 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, number AIAA-2006-7117, Portsmouth, VA, September 6–8 2006. 192, 202, 203

[23] M. S. Eldred, A. A. Giunta, and S. S. Collis. Second-order corrections for surrogate-based optimization with model hierarchies. In *Proceedings of the 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Albany, NY,, Aug. 30–Sept. 1, 2004. AIAA Paper 2004-4457. 1063, 1191, 1192, 1210, 1211

[24] G. M. Fadel, M. F. Riley, and J.-F. M. Barthelemy. Two point exponential approximation method for structural optimization. *Structural Optimization*, 2(2):117–124, 1990. 1203

[25] R. Fourer, D. M. Gay, and B. W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming, 2nd ed.* Duxbury Press/Brooks/Cole Publishing Co., Pacific Grove, CA, 2003. For small examples, e.g., at most 300 variables, a student version of AMPL suffices; see `http://www.ampl.com/DOWNLOADS`. 1341

[26] J. H. Friedman. Multivariate adaptive regression splines. *Annals of Statistics*, 19(1):1–141, March 1991. 1165

[27] J. Gablonsky. Direct version 2.0 userguide technical report. Technical Report CRSC-TR01-08, North Carolina State University, Center for Research in Scientific Computation, Raleigh, NC, 2001. 645

[28] D. M. Gay. Hooking your solver to AMPL. Technical Report Technical Report 97-4-06, Bell Laboratories, Murray Hill, NJ, 1997. Available online as `http://www.ampl.com/REFS/HOOKING/index.html` and `http://www.ampl.com/REFS/hooking2.pdf` and `http://www.ampl.com/REFS/hooking2.ps.gz`. 1340, 1341

[29] D. M. Gay. Specifying and reading program input with NIDR. Technical Report SAND2008-2261P, Sandia National Laboratories, 2008. Available as `http://www.sandia.gov/~dmgay/nidr08.pdf`. 17

[30] R. Ghanem and J. R. Red-Horse. Propagation of probabilistic uncertainty in complex physical systems using a stochastic finite element technique. *Physica D*, 133:137–144, 1999. 108, 112

[31] R. G. Ghanem and P. D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer-Verlag, New York, 1991. 108, 112

[32] W.R. Gilks, S. Richardson, and D.J. Spiegelhalter. *Markov chain Monte Carlo in practice*. Chapman & Hall, 1998. 120

[33] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright. User's guide for NPSOL (Version 4.0): A Fortran package for nonlinear programming. Technical Report TR SOL-86-2, System Optimization Laboratory, Stanford University, Stanford, CA, 1986. 133

[34] P. E. Gill, W. Murray, and M. H. Wright. *Practical Optimization*. Academic Press, San Diego, CA, 1981. 115

[35] A. A. Giunta. Use of data sampling, surrogate models, and numerical optimization in engineering design. In *Proc. 40th AIAA Aerospace Science Meeting and Exhibit*, number AIAA-2002-0538, Reno, NV, January 2002. 1193, 1211

[36] A. A. Giunta, L. P. Swiler, S. L Brown, M. S. Eldred, M. D. Richards, and E. C. Cyr. The surfpack software library for surrogate modeling of sparse, irregularly spaced multidimensional data. In *Proceedings of the 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, number AIAA-2006-7049, Portsmouth, VA, 2006. 1149

[37] G. A. Gray and T. G. Kolda. Algorithm 856: APPSPACK 4.0: Asynchronous parallel pattern search for derivative-free optimization. *ACM Transactions on Mathematical Software*, 32(3):485–507, September 2006. 132, 451

[38] M. Gunburger and J. Burkardt. Uniformity measures for point samples in hypercubes, 2004. Available on John Burkardt's web site: `http://www.csit.fsu.edu/~burkardt/`. 1026, 1034, 1114

[39] Heikki Haario, Marko Laine, Antonietta Mira, and Eero Saksman. Dram: Efficient adaptive mcmc. *Statistics and Computing*, 16:339–354, 2006. 120, 977

[40] R. T. Haftka. Combining global and local approximations. *AIAA Journal*, 29(9):1523–1525, 1991. 1191, 1210

[41] R. T. Haftka and Z. Gurdal. *Elements of Structural Optimization*. Kluwer, Boston, 1992. 115

[42] A. Haldar and S. Mahadevan. *Probability, Reliability, and Statistical Methods in Engineering Design*. Wiley, New York, 2000. 108, 1061, 1309

[43] J. H. Halton. On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numerische Mathematik*, 2:84–90, 1960. 1038, 1111, 1112

[44] J. H. Halton and G. B. Smith. Algorithm 247: Radical-inverse quasi-random point sequence. *Communications of the ACM*, 7:701–702, 1964. 1038, 1111, 1112

[45] W. E. Hart, A. A. Giunta, A. G. Salinger, and B. G. van Bloemen Waanders. An overview of the adaptive pattern search algorithm and its application to engineering optimization problems. In *Proceedings of the McMaster Optimization Conference: Theory and Applications*, Hamilton, Ontario, Canada, 2001. 567, 574

[46] J. C. Helton and F. J. Davis. Sampling-based methods for uncertainty and sensitivity analysis. Technical Report SAND99-2240, Sandia National Laboratories, Albuquerque, NM, 2000. 107, 110

[47] J. C. Helton and W. L. Oberkampf. Special issue of reliability engineering and system safety: Issue on alternative representations of epistemic uncertainty, Jul–Sep 2004. 114

[48] D. Higdon, J. Gattiker, B. Williams, and M. Rightley. Computer model calibration using high-dimensional output. *Journal of the American Statistical Association*, 103(482):570–583, 2008. 121, 986

[49] N. J. Higham. The numerical stability of barycentric lagrange interpolation. *IMA Journal of Numerical Analysis*, 24(4):547–556, 2004. 112

[50] M. Hohenbichler and R. Rackwitz. Improvement of second-order reliability estimates by importance sampling. *J. Eng. Mech., ASCE*, 114(12):2195–2199, 1988. 1063

[51] H.P. Hong. Simple approximations for improving second-order reliability estimates. *J. Eng. Mech., ASCE*, 125(5):592–595, 1999. 1063

[52] R. L. Iman and W. J. Conover. A distribution-free approach to inducing rank correlation among input variables. *Communications in Statistics: Simulation and Computation*, B11(3):311–334, 1982. 1309

[53] R. L. Iman and M. J Shortencarier. A Fortran 77 program and user's guide for the generation of latin hypercube samples for use with computer models. Technical Report NUREG/CR-3624, SAND83-2365, Sandia National Laboratories, Albuquerque, NM, 1984. 107, 110

[54] E. T. Jaynes and G. Larry. Bretthorst. *Probability theory : the logic of science*. Cambridge University Press, Cambridge, UK; New York, NY, 2003. 120

[55] D. Jones, M. Schonlau, and W. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 1998. 652

[56] M. C. Kennedy and A. O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society*, 63:425–464, 2001. 120, 121, 986

[57] W. A. Klimke. *Uncertainty Modeling using Fuzzy Arithmetic and Sparse Grids*. PhD thesis, Universität Stuttgart, Stuttgart, Germany, 2005. 112

[58] R. M. Lewis and S. N. Nash. A multigrid approach to the optimization of systems governed by differential equations. In *Proceedings of the 8th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, number AIAA-2000-4890, Long Beach, CA, Sep 2000. 1191, 1192, 1210

[59] J. M. McFarland. *Uncertainty Analysis for Computer Simulations through Validation and Calibration*. PhD thesis, Vanderbilt University, Nashville, Tennesssee, 2008. available for download at http://etd.library.vanderbilt.edu/ETD-db/available/etd-03282008-125137/. 655, 909, 918, 942, 951, 967, 1007, 1091, 1153

[60] M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245, 1979. 107, 110

[61] J. C. Meza, R. A. Oliva, P. D. Hough, and P. J. Williams. OPT++: an object oriented toolkit for nonlinear optimization. *ACM Transactions on Mathematical Software*, 33(2), 2007. 133

[62] J. More and D. Thuente. Line search algorithms with guaranteed sufficient decrease. *ACM Transactions on Mathematical Software*, 20(3):286–307, 1994. 373, 374, 375, 390, 391, 392, 393, 408, 409, 410, 425, 426, 427

[63] M. D. Morris. Factorial sampling plans for preliminary computational experiments. *Technometrics*, 33(2):161–174, 1991. 1041

[64] R. H. Myers and D. C. Montgomery. *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*. John Wiley & Sons, Inc., New York, 1995. 1178

[65] A. Nealen. A short-as-possible introduction to the least squares, weighted least squares, and moving least squares methods for scattered data approximation and interpolation. Technical report, Discrete Geometric Modeling Group, Technishe Universitaet, Berlin, Germany, 2004. 1169

[66] J. Nocedal and Wright S. J. *Numerical Optimization*. Springer Series in Operations Research. Springer, New York, 1999. 115

[67] W. .L. Oberkampf and J. C. Helton. Evidence theory for engineering applications. Technical Report SAND2003-3559P, Sandia National Laboratories, Albuquerque, NM, 2003. 114

[68] M. J. L. Orr. Introduction to radial basis function networks. Technical report, University of Edinburgh, Edinburgh, Scotland, 1996. 1174

[69] V. M. Pérez, J. E. Renaud, and L. T. Watson. An interior-point sequetial approximation optimization methodology. *Structural and Multidisciplinary Optimization*, 27(5):360–370, July 2004. 211

[70] T. D. Plantenga. HOPSPACK 2.0 user manual. Technical Report SAND2009-6265, Sandia National Laboratories, 2009. 132

[71] E. Prudencio and S. H. Cheung. Parallel adaptive multilevel sampling algorithms for the bayesian analysis of mathematical models. *International Journal for Uncertainty Quantification*, 2:215–237, 2012. 977, 983

[72] D. G. Robinson and C. Atcitty. Comparison of quasi- and pseudo-monte carlo sampling for reliability and uncertainty analysis. In *Proceedings of the AIAA Probabilistic Methods Conference*, number AIAA99-1589, St. Louis, MO, 1999. 1118

[73] M. Rosenblatt. Remarks on a multivariate transformation. *Annals of Mathematical Statistics*, 23(3):470–472, 1952. 111

[74] A. Saltelli, S. Tarantola, F. Campolongo, and M. Ratto. *Sensitivity Analysis in Practice: A Guide to Assessing Scientific Models.* John Wiley & Sons, 2004. 723, 769, 770, 802, 1027, 1034, 1035, 1041, 1115

[75] K. Schittkowski. NLPQLP: A fortran implementation of a sequential quadratic programming algorithm with distributed and non-monotone line search – user's guide. Technical report, Department of Mathematics, University of Bayreuth, Bayreuth, Germany, 2004. 133

[76] G. D. Sjaardema. APREPRO: An algebraic preprocessor for parameterizing finite element analyses. Technical Report SAND92-2291, Sandia National Laboratories, Albuquerque, NM, 1992. 1349, 1360

[77] R. Srinivasan. *Importance Sampling*. Springer-Verlag, 2002. 819

[78] A. Stroud. *Approximate Calculation of Multiple Integrals*. Prentice Hall, 1971. 681

[79] L. P. Swiler and N. J. West. Importance sampling: Promises and limitations. In *Proceedings of the 12th AIAA Non-Deterministic Approaches Conference*, number AIAA-2010-2850, 2010. 819

[80] G. Tang, L. P. Swiler, and M. S Eldred. Using stochastic expansion methods in evidence theory for mixed aleatory-epistemic uncertainty quantification. In *Proceedings of the 51st AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference (12th AIAA Non-Deterministic Approaches conference)*, Orlando, FL, April 12-15, 2010. AIAA Paper 2010-XXXX. 108

[81] R. A. Tapia and M. Argaez. Global convergence of a primal-dual interior-point newton method for nonlinear programming using a modified augmented lagrangian function. (In Preparation). 377, 395, 412, 429

[82] C. H. Tong. The PSUADE software library. Web site, 2005. http://www.llnl.gov/CASC/uncertainty_quantification/#psuade. 134

[83] R. J. Vanderbei and D. F. Shanno. An interior-point algorithm for nonconvex nonlinear programming. *Computational Optimization and Applications*, 13:231–259, 1999. 378, 395, 412, 429

[84] G. N. Vanderplaats. CONMIN – a FORTRAN program for constrained function minimization. Technical Report TM X-62282, NASA, 1973. See also Addendum to Technical Memorandum, 1978. 131, 296

[85] G. N. Vanderplaats. *Numerical Optimization Techniques for Engineering Design: With Applications*. McGraw-Hill, New York, 1984. 115

[86] Vanderplaats Research and Development, Inc., Colorado Springs, CO. *DOT Users Manual, Version 4.20*, 1995. 131, 263

[87] J. A. Vrugt, C. J. F. ter Braak, C. G. H. Diks, B. A. Robinson, J. M. Hyman, and D. Higdon. Accelerating markov chain monte carlo simulation by self-adaptive differential evolution with randomized subspace sampling. *International Journal of Nonlinear Scientific Numerical Simulation*, 10(3), 2009. 121, 962, 1000

[88] V. G. Weirs, J. R. Kamm, L. P. Swiler, M. Ratto, S. Tarantola, B. M. Adams, W. J. Rider, and M. S Eldred. Sensitivity analysis techniques applied to a system of hyperbolic conservation laws. *Reliability Engineering and System Safety*, 107:157–170, 2012. 723, 770, 802, 1027, 1035, 1115

[89] S. J. Wright. *Primal-Dual Interior-Point Methods*. SIAM, 1997. 379, 396, 413, 430

[90] G. D. Wyss and K. H. Jorgensen. A user's guide to LHS: Sandia's Latin hypercube sampling software. Technical Report SAND98-0210, Sandia National Laboratories, Albuquerque, NM, 1998. 1252, 1255

[91] D. Xiu. Numerical integration formulas of degree two. *Applied Numerical Mathematics*, 58:1515–1520, 2008. 681

[92] S. Xu and R. V. Grandhi. Effective two-point function approximation for design optimization. *AIAA J.*, 36(12):2269–2275, 1998. 1063, 1202, 1203

[93] D. C. Zimmerman. Genetic algorithms for navigating expensive and complex design spaces, September 1996. Final Report for Sandia National Laboratories contract AO-7736 CA 02. 1171