

Strategies for Multifidelity Optimization with Variable Dimensional Hierarchical Models

T. D. Robinson^{||*}; M. S. Eldred^{**†}; K. E. Willcox^{||‡}; R. Haimes^{||§}

^{||}*Aerospace Computational Design Laboratory
Massachusetts Institute of Technology, Cambridge, MA, 02139*

^{**}*Optimization and Uncertainty Estimation Department
Sandia National Laboratories,[¶] Albuquerque, NM 87185*

Surrogate-based-optimization methods are increasingly used to minimize expensive high-fidelity models and therefore reduce the computational cost. The methods are useful in problems for which two models of the same physical system exist: a high-fidelity model which is accurate and expensive, and a low-fidelity model which is cheaper but less accurate. A number of model management techniques have been developed and shown to work well for the case in which both models are defined over the same design space. However, many systems exist with variable fidelity models for which the design variables are defined over different spaces, and a mapping is required between the spaces. Three mapping methods are presented: space mapping, corrected space mapping, and proper orthogonal decomposition (POD) mapping. The methods are used within a trust-region model-management framework. They are applied to three example problems: a fixed-dimension analytic problem, a variable-dimension analytic problem, and a variable-dimension airfoil design problem. Both corrected space mapping and POD mapping are provably convergent to an optimum of the high-fidelity problem, and for the airfoil design problem considered, the POD mapping yields significant computational savings over direct optimization on the high-fidelity model.

I. Introduction

DESIGNERS of engineering systems are increasingly using numerical optimization methods. These methods enable many design options to be explored systematically, thus offering the potential of increased product performance and profitability. While optimization methods have predominantly been used in the preliminary and detailed design phases, there is an increasing need to include higher-fidelity models, such as simulation-based physics models, earlier in the design process. However, these models tend to have high computational cost even for the evaluation of a single design. Therefore, optimization, which requires a large number of evaluations, can be prohibitively expensive. One method to reduce the cost of optimization is to use surrogate-based-optimization (SBO) methods. This paper presents new methodology for extending SBO methods to handle a broader hierarchy of design models that may encountered in conceptual and preliminary design.

Low-fidelity models can be roughly divided into three categories: data fits, typically using interpolation or regression of the high-fidelity model evaluated at one or more sample points;¹ reduced-order models, derived using techniques such as modal analysis^{2,3} and proper orthogonal decomposition (POD);^{4,5} and hierarchical models,

*Graduate Student, Student Member AIAA, robinst@mit.edu

†Principal Member of Technical Staff, Associate Fellow AIAA

‡Associate Professor of Aeronautics and Astronautics, Senior Member AIAA

§Principal Research Engineer, Member AIAA

¶Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed-Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

also called multifidelity, variable-fidelity, or variable-complexity models. In this last case, a physics-based model of lower accuracy and reduced computational cost is used in conjunction with the high-fidelity model. The multifidelity case can be further divided based on the means by which the fidelity is reduced. The low-fidelity surrogate can be the same as the high-fidelity, but converged to a higher residual tolerance; for finite element models, it can use a lower basis function order than the high-fidelity model; it can be the same model on a coarser grid;^{6,7} or it can use a simpler engineering model that neglects some physics modeled by the high-fidelity method.⁸

In the latter two cases, sometimes the lower-fidelity model requires a different number of design variables than the higher-fidelity model. In the variable-grid-size case, the design variables can be discretized in the same manner as the grid, leading to more design variables on the finer grid. In the neglected-physics case, the two models often use entirely different sets of design variables. For example, in aircraft design a range of aerodynamic models can be used, from high-fidelity computational fluid dynamics (CFD) models to lower fidelity predictions using classical aerodynamic theory or vortex-lattice models. The corresponding parameterization of the design vector also varies. For example, a detailed description of aircraft geometric parameters such as airfoil cross-sectional shapes may be used in conjunction with a CFD model, while higher level design variables, such as wing area, aspect ratio and wing sweep, are appropriate for use with lower fidelity models.⁹

Surrogate models can be rigorously incorporated into design optimization through the use of a formal model management strategy. One such strategy is a trust-region model-management (TRMM) framework.¹⁰ By using corrections to ensure that the surrogate model is at least first-order accurate at the center of the trust region, this method is provably convergent to a local minimum of the high-fidelity function. The TRMM framework is widely used, having been adapted for multi-objective optimization¹¹ and multidisciplinary optimization,¹² in cases when the design vector is the same between the high-fidelity and low-fidelity models.

A general unconstrained design problem can be posed using the following nonlinear optimization formulation:

$$\text{minimize } f(\mathbf{x}), \tag{1}$$

where $f \in \mathbb{R}$ represents the scalar objective to be minimized and $\mathbf{x} \in \mathbb{R}^n$ is the vector of n design variables that describe the design. A low-fidelity model of $f(\mathbf{x})$ is denoted $g(\tilde{\mathbf{x}})$, where $\tilde{\mathbf{x}} \in \mathbb{R}^{\tilde{n}}$ is the low-fidelity design vector of dimension \tilde{n} . Many authors use the terms “surrogate” and “low-fidelity model” interchangeably. However, in this work, a surrogate model $\tilde{f}(\mathbf{x})$ is a model defined over the same design variable space \mathbf{x} as the high-fidelity model, and that meets certain consistency conditions with respect to the high-fidelity model.¹³ A low-fidelity model $g(\tilde{\mathbf{x}})$ is a model of the same physical system, but can be defined over a different design space, $\tilde{\mathbf{x}} \neq \mathbf{x}$.

SBO methods have until now been applicable only to models in which both the high-fidelity model $f(\mathbf{x})$ and the low-fidelity model $g(\tilde{\mathbf{x}})$ are defined over the same design variable space, $\mathbf{x} = \tilde{\mathbf{x}}$. The work presented in this paper provides new methodology that fills the need to extend these existing variable-fidelity optimization methods to the case when the low-fidelity model uses a different set of design variables $\tilde{\mathbf{x}} \neq \mathbf{x}$. This methodology allows variable-fidelity optimization methods to be applied more broadly across a hierarchy of design models.

This paper begins by presenting some necessary background for the work: trust-region model management methods and the corrections necessary to make those methods convergent. Three methods for mapping between variable-dimensional spaces are introduced: space mapping, corrected space mapping, and POD mapping. The mapping methods are implemented in a multifidelity optimization framework that uses TRMM. The multifidelity methods are experimentally compared with one another and with optimization applied directly to the high-fidelity problem for three examples. Finally, the relative advantages and disadvantages of the methods are discussed and recommendations for their use are presented.

II. Background

A. Trust-region model management

Surrogate-based optimization methods can be implemented within a TRMM framework. At each iteration, a surrogate for the high-fidelity model is optimized within a trust region. For the next iteration, the size of the trust region grows if the surrogate is accurately predicting the performance of the high-fidelity model, or shrinks if the surrogate is inaccurate. The TRMM framework used in this work is adapted from Ref. 10. TRMM can be applied to both unconstrained and constrained optimization problems, and future work will apply constrained TRMM

to problems with variable design space models. The discussion below, however, considers only an unconstrained problem.

In order to solve the general nonlinear programming problem given in (1), the method solves a sequence of trust region optimization subproblems. In this approach, the design variables of each subproblem are the same as those of the original high-fidelity problem. The surrogate model can change from iteration to iteration; normally it is the same low-fidelity model incorporating different corrections, as described in Section II. B. The k^{th} subproblem takes the form

$$\begin{aligned} & \text{minimize} && \tilde{f}^k(\mathbf{x}) \\ & \text{subject to} && \|\mathbf{x} - \mathbf{x}_c^k\|_\infty \leq \Delta^k, \end{aligned}$$

where \tilde{f}^k denotes the k^{th} surrogate model, \mathbf{x}_c^k is the solution to the previous subproblem and the center point of the trust region for subproblem k , Δ^k is the size of the trust region for subproblem k , and the initial trust-region size Δ^0 is user-selected. The solution to the k^{th} subproblem is denoted \mathbf{x}_*^k . After each of the k iterations in the TRMM strategy, the predicted step is validated by computing the trust region ratio ρ^k as

$$\rho^k = \frac{f(\mathbf{x}_c^k) - f(\mathbf{x}_*^k)}{\tilde{f}(\mathbf{x}_c^k) - \tilde{f}(\mathbf{x}_*^k)}, \quad (2)$$

which is the ratio of the actual improvement in the objective function to the improvement predicted by optimization on the surrogate model. This ratio measures the performance of the corrected low-fidelity model in finding new iterates that improve the high-fidelity objective. The value for ρ then defines the step acceptance and the next trust region size Δ^k using the following logic:

1. $\rho^k \leq 0$: The surrogates are inaccurate. Reject the step and shrink the trust region by half to improve surrogate accuracy.
2. $0 < \rho^k \leq r_1$: The surrogates are marginally accurate. Accept the step but shrink the trust region size by half.
3. $r_1 < \rho^k < r_2$: The surrogates are moderately accurate. Accept the step and maintain the current trust region size.
4. $r_2 \leq \rho^k$: The surrogates are accurate. Accept the step and increase the trust region size by a factor of two.

This work uses $r_1 = 10^{-5}$ and $r_2 = 0.8$. The algorithm calls the high-fidelity analysis method once per iteration and the surrogate analysis method many times per iteration. The TRMM method is provably convergent to a local minimum of the high-fidelity function if the surrogate is at least first-order accurate at the center of the trust region.¹⁴ First-order accuracy can be guaranteed through the use of corrections, as described in the following subsection.

B. Corrections

An appropriate surrogate model $\tilde{f}(\mathbf{x})$, is one that is at least first-order accurate at the center of the trust region. Given any low-fidelity model $g(\mathbf{x})$ that approximates $f(\mathbf{x})$, a consistent surrogate model can be derived through the use of corrections.

A variety of relationships between the high- and low-fidelity models can be used in deriving correction approaches. The primary two of interest are

$$A(\mathbf{x}) = f(\mathbf{x}) - g(\mathbf{x}) \quad (3)$$

$$B(\mathbf{x}) = \frac{f(\mathbf{x})}{g(\mathbf{x})}, \quad (4)$$

which correspond to the (exact) additive correction $f(\mathbf{x}) = g(\mathbf{x}) + A(\mathbf{x})$ and to the (exact) multiplicative correction $f(\mathbf{x}) = g(\mathbf{x}) B(\mathbf{x})$. Approximations to the exact correction functions $A(\mathbf{x})$ and $B(\mathbf{x})$ are denoted $\alpha(\mathbf{x})$ and $\beta(\mathbf{x})$ respectively.

The second-order Taylor series expansions of A and B around \mathbf{x}_c are:

$$\alpha(\mathbf{x}) = A(\mathbf{x}_c) + \nabla A(\mathbf{x}_c)^T (\mathbf{x} - \mathbf{x}_c) + \frac{1}{2} (\mathbf{x} - \mathbf{x}_c)^T \nabla^2 A(\mathbf{x}_c) (\mathbf{x} - \mathbf{x}_c) \quad (5)$$

$$\beta(\mathbf{x}) = B(\mathbf{x}_c) + \nabla B(\mathbf{x}_c)^T (\mathbf{x} - \mathbf{x}_c) + \frac{1}{2} (\mathbf{x} - \mathbf{x}_c)^T \nabla^2 B(\mathbf{x}_c) (\mathbf{x} - \mathbf{x}_c) \quad (6)$$

where, by differentiating equation (3),

$$A(\mathbf{x}_c) = f(\mathbf{x}_c) - g(\mathbf{x}_c) \quad (7)$$

$$\nabla A(\mathbf{x}_c) = \nabla f(\mathbf{x}_c) - \nabla g(\mathbf{x}_c) \quad (8)$$

$$\nabla^2 A(\mathbf{x}_c) = \nabla^2 f(\mathbf{x}_c) - \nabla^2 g(\mathbf{x}_c). \quad (9)$$

Similarly, by differentiating equation (4),

$$B(\mathbf{x}_c) = \frac{f(\mathbf{x}_c)}{g(\mathbf{x}_c)} \quad (10)$$

$$\nabla B(\mathbf{x}_c) = \frac{1}{g(\mathbf{x}_c)} \nabla f(\mathbf{x}_c) - \frac{f(\mathbf{x}_c)}{[g(\mathbf{x}_c)]^2} \nabla g(\mathbf{x}_c) \quad (11)$$

$$\begin{aligned} \nabla^2 B(\mathbf{x}_c) = & \frac{1}{g(\mathbf{x}_c)} \nabla^2 f(\mathbf{x}_c) - \frac{f(\mathbf{x}_c)}{[g(\mathbf{x}_c)]^2} \nabla^2 [g(\mathbf{x}_c)] + \frac{2f(\mathbf{x}_c)}{[g(\mathbf{x}_c)]^3} \nabla g(\mathbf{x}_c) \nabla g^T(\mathbf{x}_c) - \\ & \frac{1}{[g(\mathbf{x}_c)]^2} [\nabla g(\mathbf{x}_c) \nabla f^T(\mathbf{x}_c) + \nabla f(\mathbf{x}_c) \nabla g^T(\mathbf{x}_c)]. \end{aligned} \quad (12)$$

The corrected low-fidelity function in the additive case is then given by

$$\tilde{f}_\alpha(\mathbf{x}) = g(\mathbf{x}) + \alpha(\mathbf{x}) \quad (13)$$

and in the multiplicative case by

$$\tilde{f}_\beta(\mathbf{x}) = g(\mathbf{x})\beta(\mathbf{x}), \quad (14)$$

where $\alpha(\mathbf{x})$ and $\beta(\mathbf{x})$ are defined in equations (5) and (6).

These derivations have shown second-order corrections. For engineering problems, exact Hessian matrices $\nabla^2 f(\mathbf{x}_c)$ and $\nabla^2 g(\mathbf{x}_c)$ are often unavailable. They can be estimated using finite differences or approximated through quasi-Newton approximations,¹³ such as the Broyden-Fletcher-Goldfarb-Shanno (BFGS)¹⁵⁻¹⁸ update. For first-order corrections (sufficient to ensure provable convergence of the SBO method), $\nabla^2 A(\mathbf{x}_c)$ and $\nabla^2 B(\mathbf{x}_c)$ are neglected in equations (5) and (6). In this work, all low-fidelity functions are corrected to quasi-second-order accuracy using the BFGS approximation to the Hessian matrix. The additive correction was chosen over the multiplicative because it has been shown to be appropriate in a wider variety of problems.¹³

III. Mapping between variable-dimensional spaces

SBO methods have until now been applicable only to models in which both the high-fidelity model $f(\mathbf{x})$ and the low-fidelity model $g(\tilde{\mathbf{x}})$ are defined over the same space $\mathbf{x} = \tilde{\mathbf{x}}$. In order to use a low-fidelity model with a different number of design variables as the high-fidelity function to be optimized, it is necessary to find a relationship between the two sets of design vectors, that is, $\tilde{\mathbf{x}} = P(\mathbf{x})$ or $\mathbf{x} = Q(\tilde{\mathbf{x}})$. Normally, the first mapping is more useful, as it allows the optimization algorithm to work in high-fidelity space. That is, $g(P(\mathbf{x}))$ can be corrected to a surrogate for $f(\mathbf{x})$. In the second case, the optimization is performed in the low-fidelity space, and $g(\tilde{\mathbf{x}})$ is corrected to a surrogate for $f(Q(\tilde{\mathbf{x}}))$.

In some cases, the mapping can be obvious and problem-specific. For instance, if the high- and low-fidelity models are the same set of physical equations, but on a fine and coarse grid, and the design vectors in each case are geometric parameters defined on that grid, the low-fidelity design vector can be a subset of the high-fidelity design vector, or the high-fidelity design vector can be an interpolation of the low-fidelity design vector. However, in other problems, there is no obvious mathematical relationship between the design vectors. In this case, an empirical mapping is needed.

One example of a problem without an obvious mapping is the airfoil design problem described in Section IV. C. Another is the multifidelity supersonic business jet problem used by Choi, Alonso, and Kroo.⁹ That problem is to optimize the aerodynamic performance of a low-supersonic-boom aircraft. The low-fidelity model uses classical supersonic aerodynamics and vortex lattice-methods, and operates on an aircraft defined by 16 design variables: the wing area, aspect ratio, and sweep, the location of the wing root leading edge, the thickness to chord ratio at three locations on the wing, the minimum cockpit diameter, the minimum cabin diameter, and the fuselage radii at six locations. The high-fidelity model uses 126 design variables: leading and trailing edge droop, twist, and 15 camber Hicks-Henne bumps at each of 7 locations on the wing. The high-fidelity analysis uses the Euler equations and provides the gradient through the adjoint method. Choi et al. used the two models sequentially, optimizing first using the low-fidelity model, with Kriging corrections applied, and using the result of that optimization as a starting point for optimization using the high-fidelity model.

When applying mapping, the corrections must be modified to include the mapping and the Jacobians. For example, for optimization in the high-fidelity space, equation (8) becomes

$$\nabla_{\mathbf{x}} A(\mathbf{x}_c) = \nabla_{\mathbf{x}} f(\mathbf{x}_c) - \frac{\partial P^T}{\partial \mathbf{x}} \nabla_{\tilde{\mathbf{x}}} g[P(\mathbf{x}_c)] \quad (15)$$

where $\frac{\partial P}{\partial \mathbf{x}}$ is the Jacobian matrix of the mapping and $\nabla_{\mathbf{x}}$ and $\nabla_{\tilde{\mathbf{x}}}$ denote derivatives with respect to \mathbf{x} and $\tilde{\mathbf{x}}$, respectively.

In the following subsections, three methods of mapping between spaces of different dimension are presented: space mapping, corrected space mapping, and POD-based mapping.

A. Space mapping

Space mapping^{19,20} is a method of linking variable-fidelity models developed in the microwave circuit design community. In that application area, it is often appropriate to make corrections to the input of a model, rather than to its output. Space mapping has been suggested as a method for mapping between models of different dimensionality,¹⁹ although the examples available in the literature consider only the case where the design vectors have the same length. In space mapping, a particular form is assumed for the relationship P between the high- and low-fidelity design vectors. This form is described by some set of parameters, contained here in a vector \mathbf{p} , that are found by solving an optimization problem

$$\mathbf{p} = \arg \min_{\mathbf{p}} \sum_{i=1}^S (||f(\mathbf{x}^i) - g(P(\mathbf{x}^i, \mathbf{p}))||_2). \quad (16)$$

This optimization problem seeks to minimize the difference between the high-fidelity function $f(\mathbf{x})$ and the corresponding low-fidelity function $g(\tilde{\mathbf{x}}) = g(P(\mathbf{x}, \mathbf{p}))$ over a set of S sample points \mathbf{x}^i , where \mathbf{x}^i denotes the i^{th} sample (or control) point. Both the choice of sample points and the particular form of the mapping P is left to the implementation. Because the method does not ensure first-order accuracy, the proofs of convergence of trust-region methods do not extend to those methods using space mapping. However, Madsen and Søndergaard²¹ have developed a provably convergent algorithm by using a hybrid method in which the surrogate model is a convex combination of the space-mapped low-fidelity function and a Taylor series approximation to the high-fidelity function.

In the implementation employed in this paper, the sample points used in equation (16) are the previous S solutions of the trust region subproblems, \mathbf{x}_*^k , at which high-fidelity function values are already available, and a linear relationship is chosen for the mapping P :

$$\tilde{\mathbf{x}} = P(\mathbf{x}) = \mathbf{M}\mathbf{x} + \mathbf{b}, \quad (17)$$

where \mathbf{M} is a matrix with $n \times \tilde{n}$ elements and \mathbf{b} is a vector of length \tilde{n} for a total of $\tilde{n} \times (n + 1)$ space-mapping parameters. It should be noted that other forms of the mapping could also be used. The space mapping parameters must be determined at each iteration of the TRMM method by solving the optimization problem (16).

B. Corrected space mapping

Because space mapping does not provide provable convergence within a TRMM framework, but any surrogate that is first-order accurate does, one approach is to correct the space-mapping framework to at least first order. This can be done with any of the corrections in Section II. B. However, if the input parameters are first selected in order to match the output function at some number of control points, and a correction (either additive or multiplicative) is subsequently applied, it is likely that the correction will unnecessarily distort the match performed in the space-mapping step. This can be resolved by performing the space mapping and correction steps simultaneously, which is achieved by embedding the correction within the space mapping.

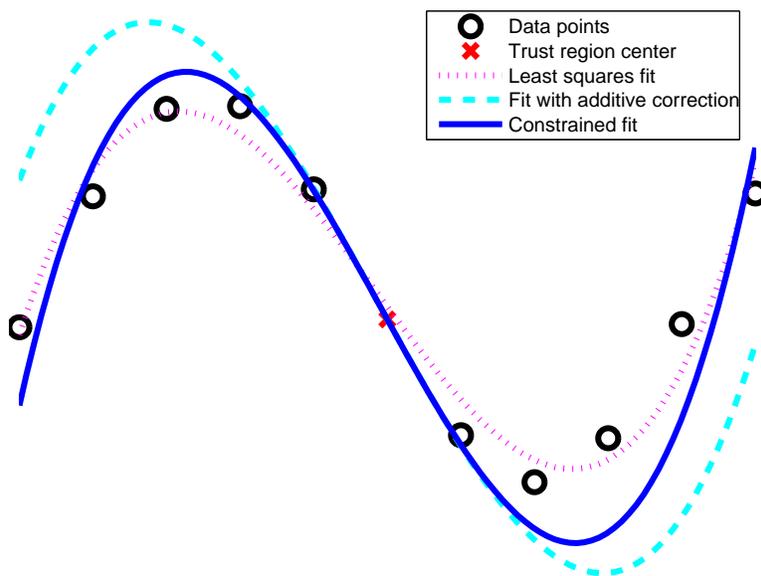


Figure 1. Demonstration of simultaneous vs. sequential data fitting and enforcement of first-order accuracy

This concept is illustrated in Figure 1. In this figure, the available data points are marked with black circles, and the center of the trust region with a red 'x'. The dotted magenta curve is a cubic function found with a least-squares fit to the available data. It provides no consistency at the trust region center. The dashed cyan curve shows the result of adding a linear additive correction to that fit, in order to enforce first-order accuracy at the center of the trust region. The local correction distorts the global data fitting. The solid blue curve is also a cubic, generated by first enforcing first-order accuracy at the center, and then performing a least-squares fit with the remaining degrees of freedom. This last curve is more globally accurate than the sequential fitting and correction steps.

Using this concept, the corrected space-mapping method performs the space mapping and correction steps simultaneously. That is, it incorporates a correction, and with the remaining degrees of freedom, performs the best match possible to the control points by varying the input mapping. The corrected surrogate to the high-fidelity function in the additive case is

$$\tilde{f}(\mathbf{x}) = g(P(\mathbf{x}, \mathbf{p})) + \alpha(\mathbf{x}, \mathbf{p}), \quad (18)$$

where α is defined in (5) and \mathbf{p} is the solution to the corrected space-mapping problem

$$\mathbf{p} = \arg \min_{\mathbf{p}} \sum_{i=1}^k \|f(\mathbf{x}^i) - \tilde{f}(\mathbf{x}^i, \mathbf{p})\|_2. \quad (19)$$

It should be noted that the correction function, $\alpha(\mathbf{x}, \mathbf{p})$, must be re-evaluated with each new value of \mathbf{p} , as the low-fidelity function value and gradient and the Jacobian of the mapping change with the space-mapping parameters. Since the resulting corrected function (18) is at least first-order accurate at the center of the trust region, the resulting trust-region optimization is provably convergent to a local minimum of the high-fidelity problem.

C. POD mapping

The third mapping methodology is based on the ‘‘gappy’’ POD method of reconstructing data sets, which is a modification of the standard POD method to handle incomplete data sets. The POD method, which is also known as principal components analysis or Karhunen-Loève expansions, yields a set of basis vectors that provides the least-squares optimal representation of a given data set.

The method of snapshots, developed by Sirovich,⁵ derives a set of POD basis vectors empirically. In this method, a set of q snapshots $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^q$, or column vectors describing different states of a system, is computed. The snapshot matrix \mathbf{X} is then formed as

$$\mathbf{X} = \begin{bmatrix} [\mathbf{x}^1 - \bar{\mathbf{x}}] & [\mathbf{x}^2 - \bar{\mathbf{x}}] & \cdots & [\mathbf{x}^q - \bar{\mathbf{x}}] \end{bmatrix}, \quad (20)$$

where $\bar{\mathbf{x}}$ is the mean of the snapshots and the i^{th} column of \mathbf{X} contains the i^{th} snapshot minus the mean. The POD basis vectors, ϕ^j , $j = 1, 2, \dots, q$, can then be computed as the left singular vectors of the matrix \mathbf{X} . The POD basis is chosen to comprise the r singular vectors that correspond to the largest singular values. A low-dimensional representation of a solution \mathbf{x} is then given by

$$\mathbf{x} \approx \bar{\mathbf{x}} + \sum_{j=1}^r \nu_j \phi^j, \quad (21)$$

where ν_j is the coefficient describing the contribution of the j^{th} POD mode ϕ^j to the solution \mathbf{x} .

It can be shown that, for a given basis size r , the POD basis yields the optimal representation of the given snapshot set in a least-squares sense. Defining for a snapshot \mathbf{x}^i , the 2-norm of the error in the low dimensional approximation as

$$e^i = \left\| \mathbf{x}^i - \bar{\mathbf{x}} - \sum_{j=1}^r \nu_j^i \phi^j \right\|_2, \quad (22)$$

then the total sum of squared errors over all snapshots is minimized and is given by the sum of the singular values corresponding to the omitted $q - r$ singular vectors,

$$\sum_{i=1}^q (e^i)^2 = \sum_{j=r+1}^q \sigma_j, \quad (23)$$

where σ_j is the j^{th} largest singular value of the matrix \mathbf{X} .

The gappy POD method, developed by Everson and Sirovich,²² allows one to reconstruct data from a ‘‘gappy’’ data set, that is, a set in which some of the data are unknown or missing.²³ The first step is to define a mask vector, which describes for a particular solution vector where data are available and where data are missing. For the solution \mathbf{x}^j , the corresponding mask vector \mathbf{n} is defined as follows:

$$\begin{aligned} n_i &= 0 \text{ if } x_i \text{ is missing} \\ n_i &= 1 \text{ if } x_i \text{ is known,} \end{aligned}$$

where x_i denotes the i^{th} element of the vector \mathbf{x} . Pointwise multiplication is defined as $(\mathbf{n}, \mathbf{x})_i = n_i x_i$. Then the gappy inner product is defined as $(\mathbf{u}, \mathbf{v})_n = ((\mathbf{n}, \mathbf{u}), (\mathbf{n}, \mathbf{v}))$, and the induced norm is $(\|\mathbf{v}\|_n)^2 = (\mathbf{v}, \mathbf{v})_n$.

For a vector \mathbf{u} that has some unknown components, it is assumed that the repaired vector $\hat{\mathbf{u}}$ can be represented by the expansion (21). In this representation, the POD coefficients ν_i are chosen to minimize the error between the available and reconstructed data. This error can be defined as

$$\epsilon = \|\mathbf{u} - \hat{\mathbf{u}}\|_n^2 \quad (24)$$

using the gappy norm so that only the original existing data elements in \mathbf{u} are compared. The coefficients ν_i that minimize the error ϵ can be found by differentiating (24) with respect to each of the ν_i in turn. This leads to the linear system of equations

$$\mathbf{E}\nu = \mathbf{w}, \quad (25)$$

where the ij^{th} component of \mathbf{E} is given by

$$E_{ij} = (\phi^i, \phi^j)_n \quad (26)$$

and the i^{th} component of \mathbf{w} is given by

$$w_i = (\mathbf{u}, \phi^i)_n. \quad (27)$$

Solving equation (25) for ν , the missing elements of \mathbf{u} can be obtained using the expansion (21).

The gappy POD method provides a way to map between high- and low-fidelity design space data: the high-fidelity vector is treated as the known data, and the low-fidelity as the unknown data, or vice versa. In the mapping application, the POD basis vectors must span both low- and high-fidelity design space. This is achieved by generating a set of q training pairs, for which the low- and the high-fidelity vectors describe the same physical system. These training pairs are combined in the following way to form the snapshot matrix:

$$\mathbf{X} = \begin{bmatrix} [\tilde{\mathbf{x}}_1 - \bar{\mathbf{x}}] & [\tilde{\mathbf{x}}_2 - \bar{\mathbf{x}}] & \cdots & [\tilde{\mathbf{x}}_r - \bar{\mathbf{x}}] \\ - & - & & - \\ [\mathbf{x}_1 - \bar{\mathbf{x}}] & [\mathbf{x}_2 - \bar{\mathbf{x}}] & \cdots & [\mathbf{x}_r - \bar{\mathbf{x}}] \end{bmatrix}, \quad (28)$$

where now the i^{th} column of \mathbf{X} contains both the i^{th} low- and the i^{th} high-fidelity snapshots, and $\bar{\mathbf{x}}$ denotes the mean of the low-fidelity snapshot set.

The left singular vectors of this snapshot matrix provide the corresponding POD basis vectors, which are partitioned in the same way as the snapshot vectors. Therefore, equation (21) can be decomposed into two equations

$$\mathbf{x} = \bar{\mathbf{x}} + \sum_{i=1}^q \nu_i \phi^i \quad (29)$$

$$\tilde{\mathbf{x}} = \bar{\mathbf{x}} + \sum_{i=1}^q \nu_i \tilde{\phi}^i, \quad (30)$$

where ϕ^i is the portion of the i^{th} POD basis vector corresponding to \mathbf{x} and $\tilde{\phi}^i$ is the portion corresponding to $\tilde{\mathbf{x}}$.

Using the gappy POD formulation, equation (29) can be solved in a least-squares sense in order to find the coefficients ν that best represent a given high-fidelity vector \mathbf{x} . Those coefficients can then be used in equation (30) to calculate the low-fidelity vector. Alternatively, if a mapping is desired from the low-fidelity space to the high-fidelity space, the coefficients are found from equation (30) and used in (29). Since POD mapping operates exclusively on the input variables, an additive or multiplicative output correction must be applied separately to ensure at least first-order consistency within the TRMM framework.

IV. Example problems

Three example problems are used to demonstrate the methods: an analytic problem in which the dimensions of the low- and high-fidelity models are the same, an analytic problem with variable-dimensional models, and an airfoil design problem with variable-dimensional models. The analytic problems are chosen to be Rosenbrock functions, which are commonly used challenging analytic optimization test cases. The trust-region model-management algorithm is applied with each of the three mapping methods and compared with the results from applying optimization directly on the high-fidelity model. Direct optimization of the high-fidelity problem was carried out using a quasi-Newton method with a BFGS approximation to the Hessian matrix.

A. Analytic problem: Rosenbrock function with the same dimension

For the first example, the dimension of the problem is the same in both the high-fidelity and low-fidelity models. The high-fidelity problem is a form of the Rosenbrock function

$$f(\mathbf{x}) = (x_1 - 1)^2 + 4(x_2 - x_1^2)^2, \quad (31)$$

and the low-fidelity model is a quadratic function

$$g(\tilde{\mathbf{x}}) = \tilde{x}_1^2 + \tilde{x}_2^2. \quad (32)$$

The optimal solution of the high-fidelity problem is $x_1 = 1$, $x_2 = 1$, with an optimal objective function value of zero. Since, in this case, the sum of the quadratic low-fidelity function and a quasi-second-order additive correction still provides a quadratic model, one would expect the convergence rate of the variable-fidelity TRMM framework to be the same as that of the direct high-fidelity optimization, which uses approximations based on a quasi-second-order Taylor series. The multifidelity methods are therefore not expected to provide any computational savings for this case, nor is a mapping from high- to low-fidelity space required in this simple case where $\mathbf{x} = \tilde{\mathbf{x}}$; however, application of the methods to this simple problem highlights the importance of using corrections in the variable-fidelity setting. For all presented results, an initial guess of $\mathbf{x} = (-2, -2)$ was used.

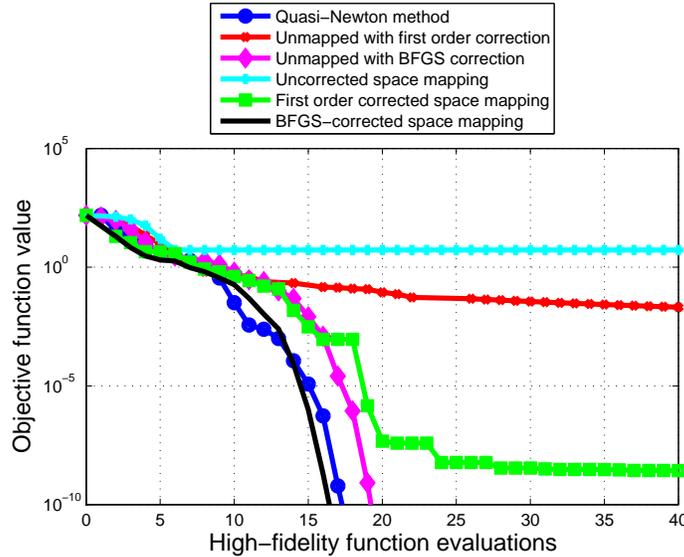


Figure 2. Multifidelity optimization, with and without space mapping, with no corrections, first-order corrections, and quasi-second-order corrections, on the two-dimensional Rosenbrock problem.

Figure 2 shows the multifidelity method using space-mapping with no corrections, first-order corrections, and BFGS corrections. At each iteration, the space-mapping parameters are selected by solving a minimization problem of the form (16). For this example, which has $n = \tilde{n} = 2$, the linear mapping requires six parameters. The space-mapping minimization is performed using the previous eight iterates as sample points. The curves labeled “unmapped” use no mapping (since a mapping is not necessary when the dimension of the two fidelities is the same) but use $\tilde{\mathbf{x}} = \mathbf{x}$ and correct the surrogate to the indicated order of consistency.

Uncorrected space mapping does not converge on this problem. With first-order corrections, the methods converge linearly. The addition of quasi-second-order corrections improves the convergence rate. As expected, the convergence rates of the BFGS corrected multifidelity methods are essentially the same as the quasi-Newton algorithm. It is also interesting to note that the convergence of the first-order corrected space mapping method is much faster than the unmapped method with a first-order correction, and almost as good as the quasi-second-order methods. One interpretation of this result is that the corrected space mapping, since it seeks a global match using information from a number of past iterations, captures some measure of curvature consistency, even though it is only formally corrected to first order.

Figure 3 shows comparisons between the multifidelity method with both POD mapping and space mapping applied. Corrected space mapping and POD mapping both use quasi-second-order corrections. The multifidelity with POD mapping shows a slight advantage on the first step, on which substantial progress is made due to the appropriate selection of the initial trust-region size, Δ^0 . Numerical experiments show that, despite some dependence in early iterations on the initial trust region size, the adaptive modification of the size of the trust regions as the TRMM algorithm progresses yields convergence rates that are essentially independent of the

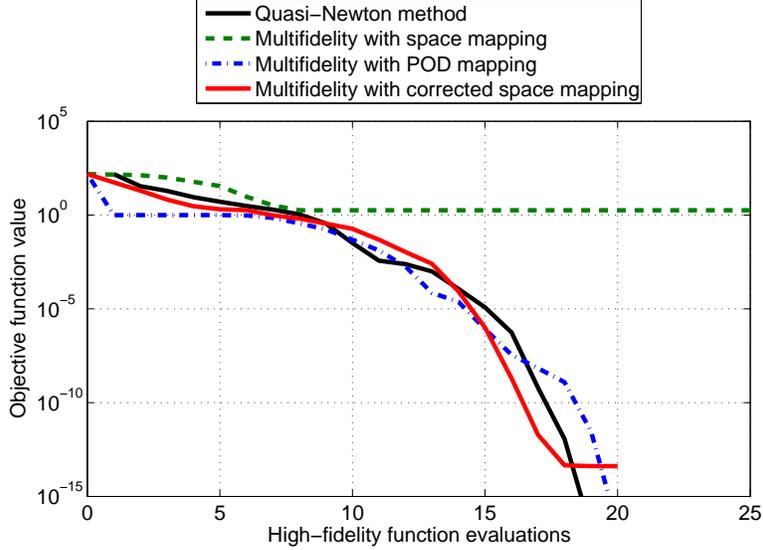


Figure 3. Comparison of optimization methods on the two-dimensional Rosenbrock problem. The corrected multifidelity methods use quasi-second-order corrections.

selection of Δ^0 . As expected for this problem, Figure 3 shows that the convergence rates for the corrected multifidelity methods and the quasi-Newton method are very similar.

B. Analytic problem: Rosenbrock function with variable dimension

The second example is a simple analytic problem for which the high- and low-fidelity design vectors have different dimension. The high-fidelity problem is a ten-dimensional extended Rosenbrock function

$$f(\mathbf{x}) = (x_1 - 1)^2 + 4 \sum_{i=2}^{10} (x_i - x_{i-1}^2)^2, \quad (33)$$

and the low-fidelity model is a two-dimensional Rosenbrock function

$$g(\tilde{\mathbf{x}}) = (\tilde{x}_1 - 1)^2 + 4(\tilde{x}_2 - \tilde{x}_1^2)^2. \quad (34)$$

For all presented results, an initial guess of $x_i = -2$, $i = 1, 2, \dots, 10$ was used.

In order to generate the snapshots required for the POD basis computation, the low-fidelity vector was varied over the grid $-3 \leq \tilde{x}_1 \leq 3$, $-3 \leq \tilde{x}_2 \leq 3$. A total of 16 snapshots were computed at equally spaced locations on this grid. The snapshots were generated using

$$\begin{aligned} x_1 &= \tilde{x}_1 \\ x_2 &= \tilde{x}_2 \\ x_i &= 0, \quad i = 3, 4, \dots, 10 \end{aligned} \quad (35)$$

for the relationship between the high- and low-fidelity design vectors. It should be noted that in this example, where the mapping is a simple linear relationship, computing all these snapshots is unnecessary and the relevant information can be captured with just two snapshots. However, the computation of the POD basis systematically identifies redundant information, yielding in this case only two non-zero singular values of the snapshot matrix.

Figure 4 shows the convergence of the multifidelity algorithm with each of the mapping methods, in addition to the standard quasi-Newton method. Because of the inherent difficulties of solving Rosenbrock optimization problems, particularly in this higher dimensional case, very slow convergence is observed for optimization applied directly to the high-fidelity problem. The quasi-Newton method is expected to eventually exhibit superlinear

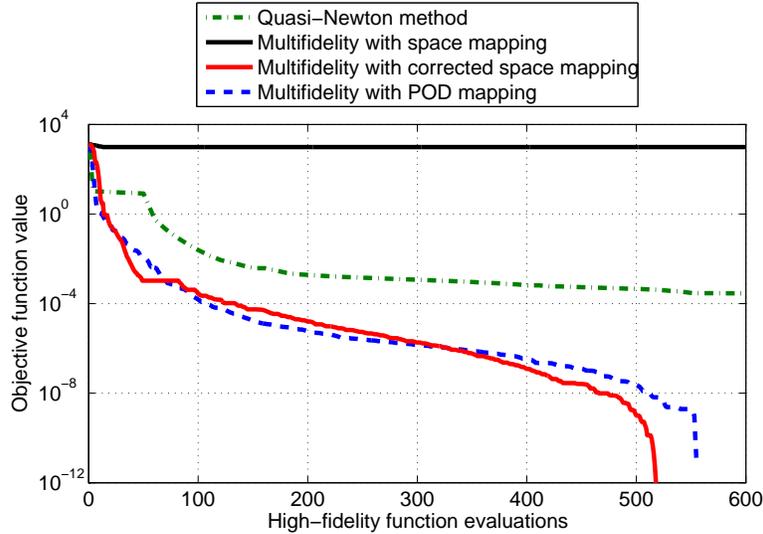


Figure 4. Comparison of optimization methods on the ten-dimensional extended Rosenbrock problem. The corrected multifidelity methods use quasi-second-order corrections.

convergence; however, it takes a very large number of iterations to do so. In contrast, the multifidelity method, with both POD mapping and corrected space mapping, converges much more quickly and is able to reach the optimal solution in less than 600 iterations. As expected, the multifidelity method with uncorrected space mapping does not converge at all, decreasing only slightly from the initial point and then failing to make additional progress. For the results shown in Figure 4, the POD mapping uses two basis vectors computed from the set of 16 snapshots.

Although the example is simple, it demonstrates the potential computational savings that can be achieved by using a low-fidelity model tailored to the problem at hand. In this case, the two-dimensional Rosenbrock function provides a much better low-fidelity approximation than the quadratic function that underlies the quasi-Newton method.

C. Airfoil design

The third problem tested was chosen to represent a realistic engineering design problem with a low-fidelity model that neglects physics contained in the high-fidelity model. This problem has a variable number of design variables due to differences in the geometric parametrization of the problem. The difference in the number of design variables is significant: the low-fidelity model has two, and the high-fidelity model has thirty-six.

The objective of this problem is to design an airfoil that matches a desired pressure distribution. The pressure distribution of the NACA 2412 airfoil was chosen as the goal. The objective function is

$$f = \int_0^1 (C_P - C_{P_{target}})^2 ds \quad (36)$$

where C_P is the coefficient of pressure and $C_{P_{target}}$ is the coefficient of pressure of the goal airfoil. The integral is over the unit chord and is approximated using trapezoidal integration. The gradient was calculated using finite differences. The initial and goal airfoils, and the corresponding coefficient of pressure distributions, are shown in Figure 5.

The low-fidelity analysis is an analytic solution to a Joukowski transform.²⁴ The Joukowski transform is a conformal map that maps the points on a unit circle to points on the surface of an airfoil. Only two variables are needed to describe a Joukowski airfoil: μ_x and μ_y , the x and y coordinates of the center of the circle used in the Joukowski transform. Figure 6 shows one Joukowski airfoil and the corresponding circle.

The high-fidelity analysis uses XFOIL.²⁵ The inviscid formulation of XFOIL uses a linear-vorticity stream function panel method. A finite trailing edge thickness is modeled with a source panel. The equations are closed

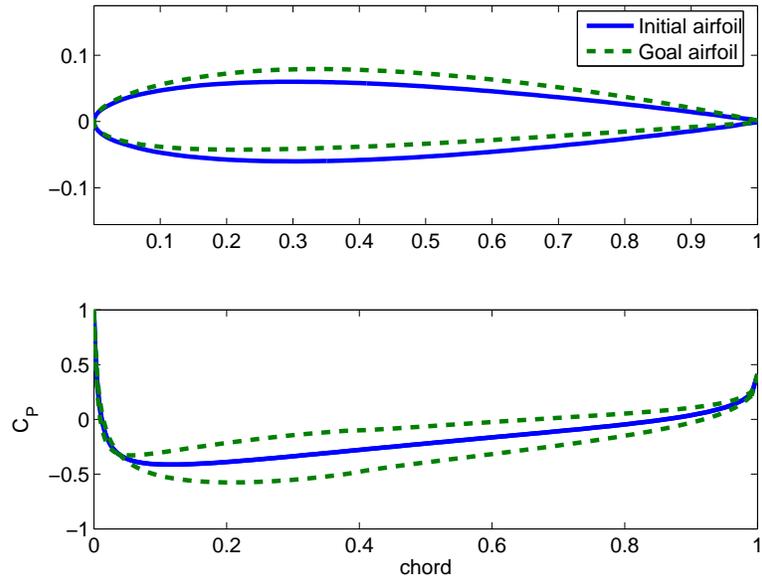


Figure 5. Initial and goal airfoils, along with their coefficient of pressure distributions.

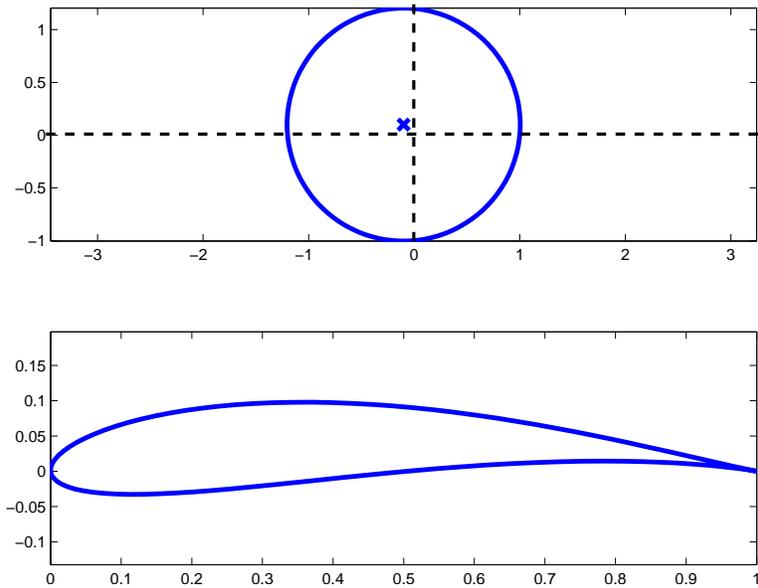


Figure 6. The unit circle in the upper plot is transformed to the airfoil in the lower plot using the Joukowski transform.

with an explicit Kutta condition. The high-fidelity geometry vector consists of the magnitudes of 36 Hicks-Henne bump parameters,²⁶ 18 of which perturb the camber of the airfoil and 18 of which perturb the thickness, as shown in Figure 7. They are evenly distributed across the airfoil at 18 control points. Therefore, this problem is a good example of a variable-fidelity problem in which the length and nature of the design vector varies significantly between the two fidelities.

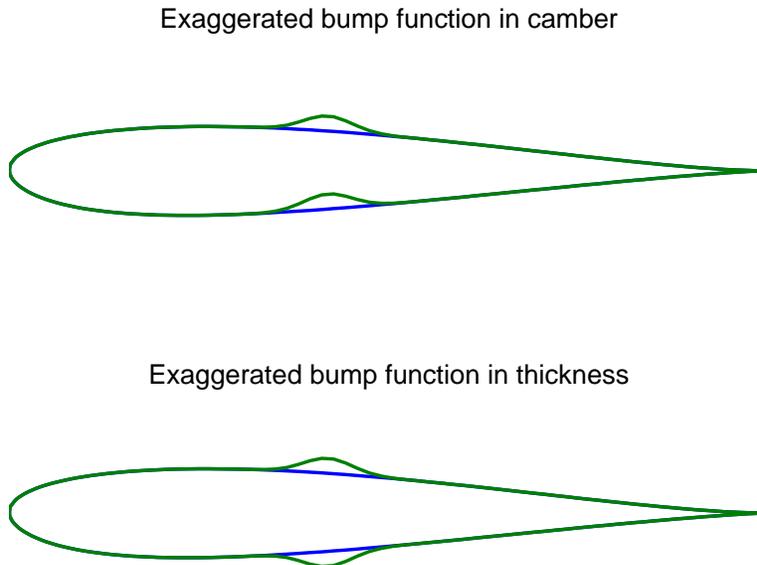


Figure 7. For the high-fidelity model, Hicks-Henne bump functions are used to parameterize the airfoil geometry.

100 snapshots were created using a grid in the Joukowski parameter space, varying μ_x from 0.01 to 0.3, and μ_y from 0 to 0.5. The corresponding high-fidelity design vectors were produced by solving an optimization problem to determine the magnitudes of the 36 Hicks-Henne bump functions that best matched the desired Joukowski airfoil. Specifically, the bump functions were chosen so as to minimize the integral over the chord length of the square of the difference between the airfoils defined in each space. The snapshots were then combined into a snapshot matrix as described in Section III. C. The mean of the snapshots was subtracted before performing the singular value decomposition.

Figure 8 shows the singular values of the POD snapshot matrix. The relative magnitudes of the singular values indicate the importance of the corresponding POD basis vectors in representing the snapshot data. The first two singular values are more than an order of magnitude larger than the third; therefore, only the first two basis vectors were used in the POD mapping.

Figures 9 – 11 show the behavior of the objective function for some selected cuts in the high-fidelity design space. These cuts correspond to an iterate near the end of the optimization process, that is, for an airfoil that is close to the goal airfoil. Figures 9 and 10 show the variation of the objective function with the variable x_{23} , which corresponds to the magnitude of a thickness bump function located approximately a third of the chord behind the leading edge, near the thickest part of the airfoil. Figure 9 shows that the corrected low-fidelity function matches the high-fidelity function well over a range of values. Figure 10 shows a more detailed section of the same plot – it can be seen that the high-fidelity function is noisy, while the corrected low-fidelity function is smooth. The corrected low-fidelity model is able to capture the second-order trend of the data even when noise is significant.

Further interrogation of the design space shows that the high-fidelity function is even noisier along variables defining the leading and trailing edges than those near the center of the airfoil. Figure 11 shows the variations in the corrected low-fidelity and high-fidelity functions with the variable x_1 , which corresponds to a bump function in camber near the leading edge. Compared with Figure 10, the oscillations in the high-fidelity function are much more significant in magnitude.

Figures 9 – 11 indicate that it is unrealistic to expect an optimization method to result in objective function values of lower than $O(10^{-4})$ for this problem. At this point, the noise in the high-fidelity function becomes significant and the progress of the optimization method is likely to be impaired due to an inability to compute

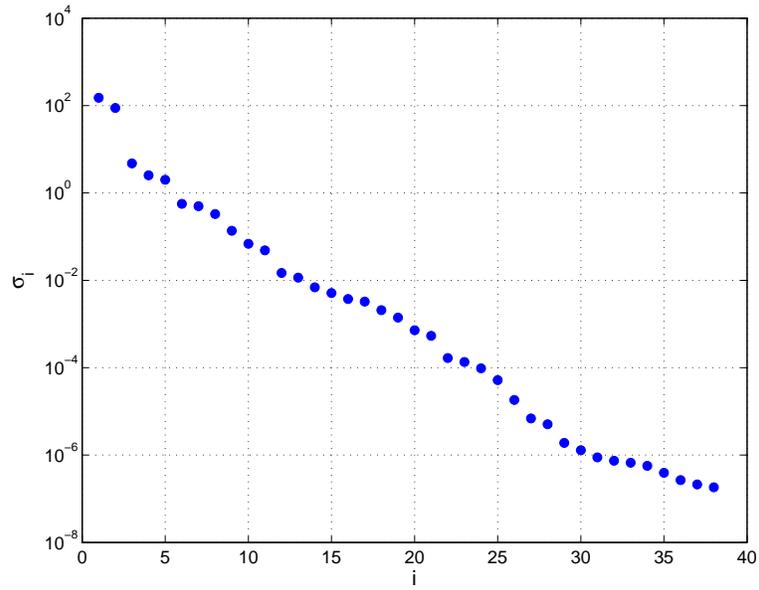


Figure 8. Singular values of the snapshot matrix for the airfoil design problem, showing that the first two are much more significant than the others.

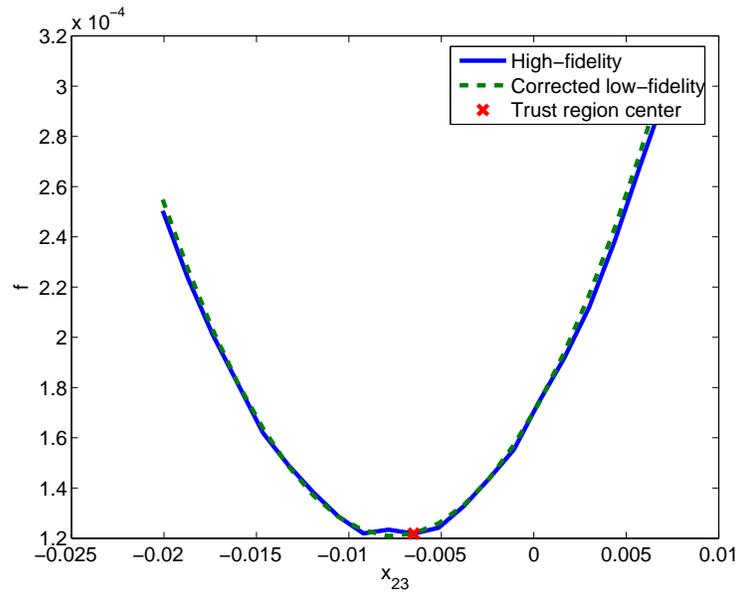


Figure 9. The high-fidelity and corrected low-fidelity functions as a function of variable x_{23} for an airfoil close to the goal airfoil.

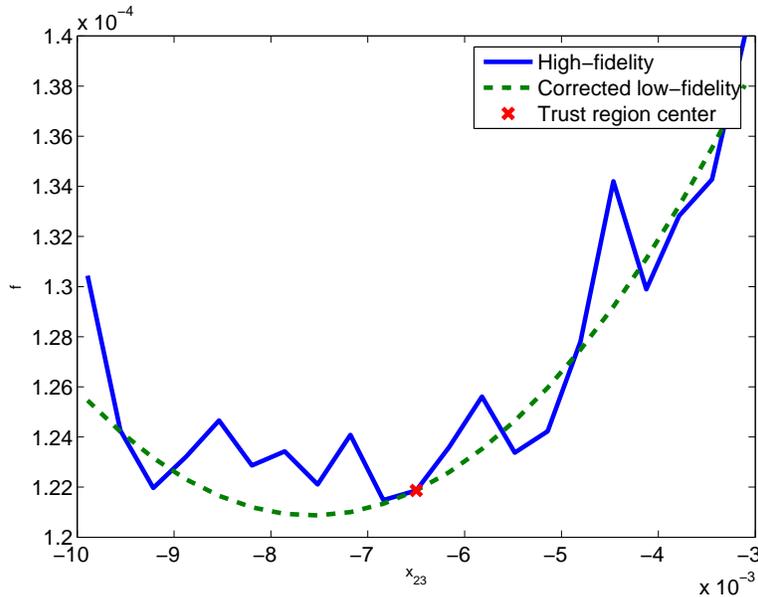


Figure 10. The high-fidelity and corrected low-fidelity functions as a function of variable x_{23} for an airfoil close to the goal airfoil, over a smaller range than Figure 9.

gradient information accurately with finite differences. A convergence criteria of 10^{-4} was therefore applied to optimization results.

Figure 12 shows the multifidelity method with each of corrected space mapping and POD mapping, along with the standard quasi-Newton method. The standard method takes 25 iterations to get to an objective function value of 10^{-4} . The multifidelity method with corrected space mapping takes 22 iterations to get to the same value, which is a small savings, and the multifidelity method with POD mapping takes 15 iterations, a savings of 40%.

Figure 13 shows the varying rates of convergence of the POD method as the number of basis vectors is varied. The curve labeled $n_{basis} = 0$ effectively uses a constant for the low-fidelity analysis. With zero basis vectors, the low-fidelity analysis always evaluates the mean airfoil. Thus, the resulting corrected surrogate function is a second-order Taylor series about the center of the trust region using the BFGS approximation to the Hessian matrix. The remaining curves show convergence with increasing numbers of basis vectors. This shows that nearly the entire mapping relationship is captured by the first basis vector and the remaining basis vectors add very little information and do not significantly improve the convergence rate.

As is generally true for gradient-based optimization of non-convex functions, the ability of the optimizer to converge to the optimal solution depends on the initial choice of design variables. The convergence rates of the multifidelity methods presented here also depend on the initial solution. In most cases, the multifidelity methods were found to yield substantial computational savings when compared with direct high-fidelity optimization; however, for some choices of initial design the improvement was not significant. The multifidelity method was applied to a number of different initial designs. Figure 14 shows four representative convergence plots. It can be seen that the POD multifidelity method is more efficient the quasi-Newton method for these cases, although for the second case there is almost no improvement.

V. Conclusion

New methodology has been presented that extends SBO methods to handle a hierarchy of design models in the case where the design vector varies between the high-fidelity and low-fidelity models. The results demonstrate the importance of corrections within a trust-region model-management framework in order to obtain provable convergence to a minimum of the high-fidelity model. For the variable-dimensional examples considered, mul-

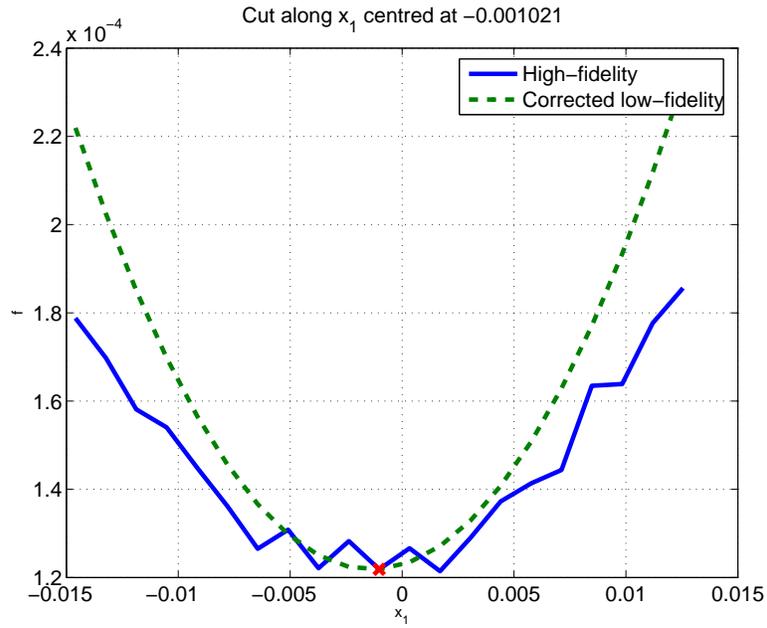


Figure 11. The high-fidelity and corrected low-fidelity functions as a function of variable x_1 for an airfoil close to the goal airfoil.

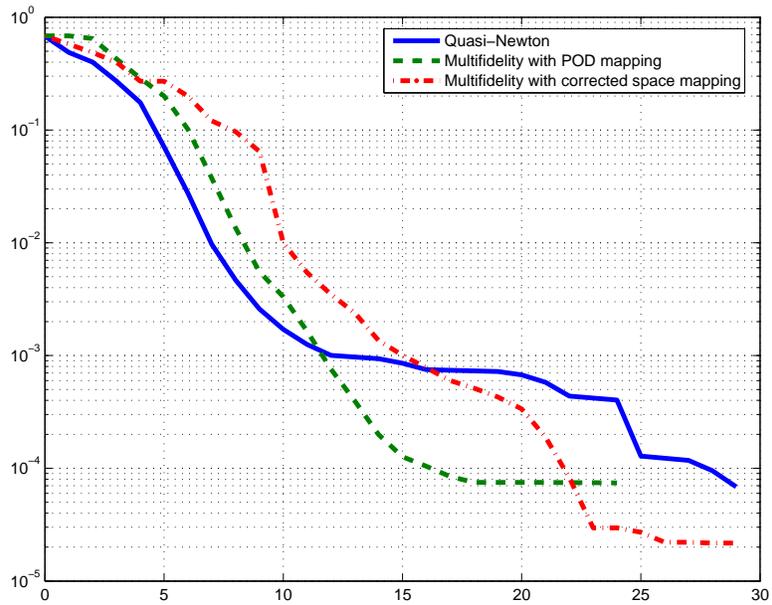


Figure 12. Convergence of quasi-Newton method, multifidelity with POD mapping, and multifidelity with corrected space mapping on the airfoil design problem.

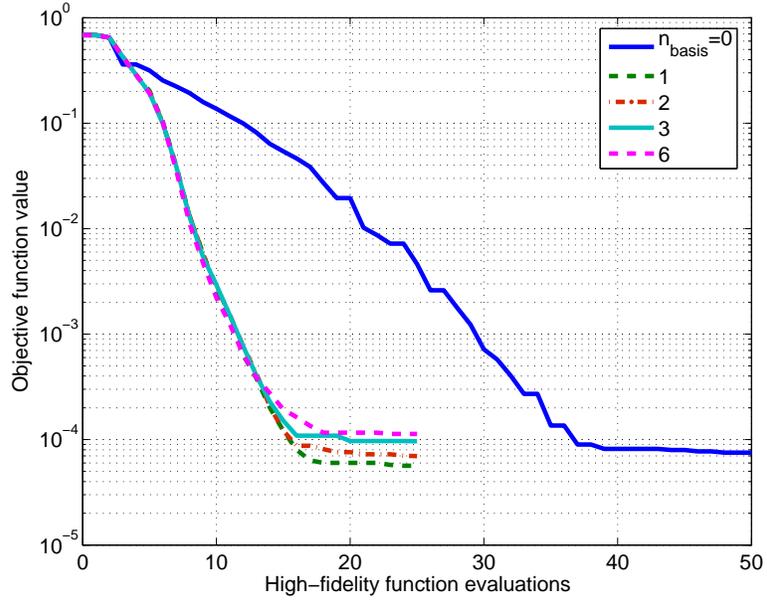


Figure 13. Convergence of multifidelity with POD mapping on airfoil design problem varying the number of basis vectors

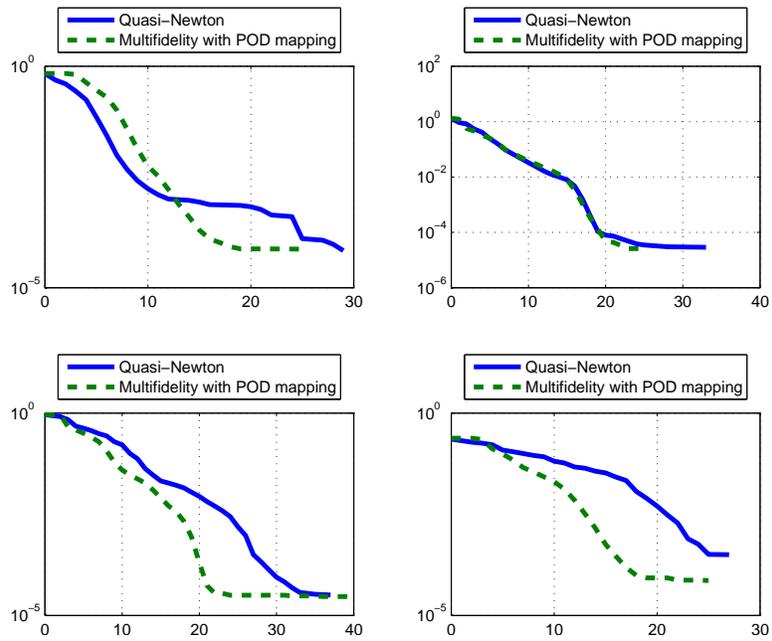


Figure 14. Convergence of multifidelity with POD mapping and a quasi-Newton method in high-fidelity space starting at four different initial airfoils

tifidelity methods that employ corrected space mapping and POD mapping were shown to provide significant computational savings over direct optimization on the high-fidelity model. POD mapping was the most efficient of the methods on the engineering design example, requiring 40% fewer high-fidelity function evaluations than the standard quasi-Newton method.

A major advantage of corrected space mapping compared to POD mapping is that it does not require training pairs, which may be non-trivial to generate. However, one must instead specify the form of the mapping. Space mapping also requires more overhead calculations on each iteration, as a small optimization problem must be solved at each iteration in order to compute the space-mapping parameters. For the POD method, the calculations for the mapping are done once, before optimization, reducing the amount of overhead on each iteration.

Ongoing work addresses the extensions of these multifidelity mapping methods to constrained problems. Several different methods for handling constraints are being investigated, including an augmented Lagrangian formulation of the problem, a method based on sequential quadratic programming, and a filter method.

Acknowledgement

This research has been supported in part by the Singapore-MIT Alliance.

References

- ¹Venter, G., Haftka, R., and J.H. Starnes, J., "Construction of Response Surface Approximations for Design Optimization," *AIAA Journal*, Vol. 36, No. 12, December 1998, pp. 2242–2249.
- ²Karpel, M., Moulin, B., and Love, M., "Modal-Based Structural Optimization with Static Aeroelastic and Stress Constraints," *Journal of Aircraft*, Vol. 34, No. 3, May-June 1997, pp. 433–440.
- ³Banerjee, J., "Exact modal analysis of an idealised whole aircraft using symbolic computation," *Aeronautical Journal*, Vol. 104, May 2000, pp. 247–255.
- ⁴Holmes, P. J., Lumley, J. L., Berkooz, G., Mattingly, J. C., and Wittenberg, R. W., "Low-dimensional models of coherent structures in turbulence," *Physics Reports*, Vol. 287, No. 4, 1997, pp. 337–384.
- ⁵Sirovich, L., "Turbulence and the Dynamics of Coherent Structures. Part 1 : Coherent Structures," *Quarterly of Applied Mathematics*, Vol. 45, No. 3, October 1987, pp. 561–571.
- ⁶Lewis, R. M. and Nash, S. G., "A Multigrid Approach to the Optimization of Systems Governed by Differential Equations," *Proceedings of the 8th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Optimization*, Long Beach, CA, September 2000, AIAA Paper 2000-4890.
- ⁷Alexandrov, N., Lewis, R., Gumbert, C., Green, L., and Newman, P., "Approximation and model management in aerodynamic optimization with variable-fidelity models," *Journal of Aircraft*, Vol. 38, No. 6, November-December 2001, pp. 1093–1101.
- ⁸Alexandrov, N. M., Nielsen, E., Lewis, R., and Anderson, W., "First-Order Model Management with Variable-Fidelity Physics Applied to Multi-Element Airfoil Optimization," *Proceedings of the 8th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, Long Beach, CA, September 2000, AIAA Paper 2000-4886.
- ⁹Choi, S., Alonso, J., Kim, S., and Kroo, I., "Two-Level Multi-Fidelity Design Optimization Studies for Supersonic Jets," *43rd AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, January 2005, AIAA-2005-531.
- ¹⁰Alexandrov, N., Dennis, J., Lewis, R., and Torczon, V., "A Trust-Region Framework for Managing the Use of Approximation Models in Optimization," *Structural and Multidisciplinary Optimization*, Vol. 15, No. 1, February 1998, pp. 16–23.
- ¹¹Yang, B., Yeun, Y., and Ruy, W., "Managing Approximation Models in Multiobjective Optimization," *Structural and Multidisciplinary Optimization*, Vol. 24, No. 2, September 2002, pp. 141–156.
- ¹²Rodriguez, J., Renaud, J., Wujek, B., and Tappeta, R., "Trust Region Model Management in Multidisciplinary Design Optimization," *Journal of Computational and Applied Mathematics*, Vol. 124, No. 1-2, December 2000, pp. 139–154.
- ¹³Eldred, M., Giunta, S., and Collis, S., "Second-order corrections for surrogate-based optimization with model hierarchies," *Proceedings of the 10th AIAA/ISSMO Multidisciplinary analysis and optimization conference*, AIAA, Albany, New York, Aug 30. - Sept. 1 2004.
- ¹⁴Alexandrov, N., Dennis, J., Lewis, R., and Torzon, V., "A Trust Region Framework for Managing the Use of Approximation Models in Optimization," NASA CR-201735, 1997.
- ¹⁵Broyden, C., "The convergence of a class of double rank minimization algorithms: Parts I and II," *Journal of the Institute of Mathematics and its Applications*, Vol. 6, 1970, pp. 76–90, 222–231.
- ¹⁶Goldfarb, D., "A family of variable metric methods derived by variational means," *Mathematics of Computation*, Vol. 24, 1970, pp. 23–26.
- ¹⁷Fletcher, R., "A New Approach to Variable Metric Algorithms," *Computer Journal*, Vol. 13, 1970, pp. 317–322.
- ¹⁸Shanno, D., "Conditioning of quasi-Newton methods for function minimization," *Mathematics of Computation*, Vol. 24, 1970, pp. 647–656.
- ¹⁹Bandler, J., Biernacki, R., Chen, S., Grobelny, P., and Hemmers, R., "Space Mapping Technique for Electromagnetic Optimization," *IEEE Transactions on Microwave Theory and Techniques*, Vol. 42, 1994, pp. 2536–2544.

²⁰Bandler, J., Cheng, Q., Dakroury, S., Mohamed, A., Bakr, M., Madsen, K., and Søndergaard, J., "Space Mapping: The State of the Art," *IEEE Transactions on Microwave Theory and Techniques*, Vol. 52, No. 1, January 2004, pp. 337–361.

²¹Madsen, K. and Søndergaard, J., "Convergence of Hybrid Space Mapping Algorithms," *Optimization and Engineering*, Vol. 5, 2004, pp. 145–156.

²²Everson, R. and Sirovich, L., "The Karhunen-Loève Procedure for Gappy Data," *Journal of the Optical Society of America*, Vol. 12, No. 8, 1995, pp. 1657–1664.

²³Bui-Thanh, T., Damodaran, M., and Willcox, K., "Aerodynamic Data Reconstruction and Inverse Design Using Proper Orthogonal Decomposition," *AIAA Journal*, Vol. 42, No. 8, August 2004, pp. 1505–1516.

²⁴Anderson, J., *Fundamentals of Aerodynamics, 2nd ed.*, McGraw-Hill, 1991.

²⁵Drela, M., "XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils," *Low Reynolds number aerodynamics : proceedings of the conference, Notre Dame, Indiana*, edited by T. Mueller, Springer-Verlag, June 1989.

²⁶Hicks, R. and Henne, P., "Wing Design by Numerical Optimization," *Journal of Aircraft*, Vol. 15, No. 7, July 1978, pp. 407–412.