

Reduced-Basis Multifidelity Approach for Efficient Parametric Study of NACA Airfoils

Ryan W. Skinner,* Alireza Doostan,† Eric L. Peters,‡ John A. Evans,§ and Kenneth E. Jansen¶
University of Colorado, Boulder, Colorado 80309

DOI: 10.2514/1.J057452

Multifidelity methods are often used to reduce the cost of parameter space exploration, with applications to design optimization and uncertainty quantification. We present a multifidelity method to construct a reduced-basis representation of spatially (or temporally) varying solution fields and use it to explore the effect of variation in the geometry and angle of attack on the pressure coefficient response of a two-dimensional NACA 4412 airfoil in steady, incompressible flow at $Re = 1.52$ Million. Two low-fidelity simulations use 1) Euler flow and 2) coarse-mesh Spalart–Allmaras Reynolds-averaged Navier–Stokes (RANS) with unresolved boundary layers. Each low-fidelity model is paired with a high-fidelity RANS model, to construct two multifidelity models that approximate the high-fidelity response associated with a set of parametric realizations. The predictive capacity and efficiency of both multifidelity models are analyzed and found to perform well for this combination of aerodynamic system and parameter space. In particular, the bifidelity model based on Euler flow predicts high-fidelity RANS results within 1.6% error at a cost that is 50 times less. Such a cost reduction demonstrates the utility of this method for industrial-scale design optimization and uncertainty quantification problems.

Nomenclature

C_p	=	pressure coefficient (this Paper's quantity of interest)
c	=	chord length, m
H	=	high-fidelity solution matrix
H^C	=	column skeleton of H
$\hat{H}$	=	approximation of H from bifidelity model
K	=	total number of low-fidelity model evaluations
L	=	low-fidelity solution matrix
L^C	=	column skeleton of L
$\hat{L}$	=	approximation of L from interpolating decomposition
m	=	maximum camber
n	=	number of points on airfoil where C_p is sampled
p	=	location of maximum camber
Re	=	Reynolds number
r	=	rank $r \ll K$, or number of high-fidelity model evaluations used to construct bifidelity model
t	=	maximum thickness
$y^{(i)}$	=	i th realization of random input parameters
$u_H(y^{(i)})$	=	output data from high-fidelity model corresponding to i th input
$u_L(y^{(i)})$	=	output data from low-fidelity model corresponding to i th input
α	=	angle of attack, deg
ϵ	=	relative error of bifidelity vs high-fidelity models
ϵ_{LvH}	=	relative error of low-fidelity vs high-fidelity models
Λ	=	coefficient matrix used in interpolating decomposition
$\ \cdot \ $	=	2-norm for vectors and matrices

I. Introduction

THE need to characterize complex aerodynamic systems and optimize their design becomes more pressing as atmospheric vehicles push performance boundaries. However, despite advances over the past five decades in turbulence modeling and computing power, accurately capturing the dynamics critical to these systems remains resource intensive. The design space exploration or uncertainty quantification (UQ) of such systems requires repeated model evaluations, which becomes computationally intractable due to their excessive cost. For problems with tractable cost, when the design or uncertainty space is characterized by a large number of parameters, a large ensemble of model evaluations is needed to thoroughly explore the high-dimensional parameter space. This may similarly lead to computational cost that exceeds available resources. Consequently, the finite time to market for most high-performance aerospace products constrains simulations' quantity and rigor, which can lead to suboptimal performance.

Multifidelity modeling** [1] is one approach for reducing the cost of such studies while maintaining reasonable accuracy. It recognizes that, in practice, engineers predicting a quantity of interest (QOI) may offload work to low-fidelity (LF) models that capture trends in a QOI predicted by high-fidelity (HF) models. LF models are typically analytic or computational and run quickly but offer only approximate predictions of the QOI. In aerodynamic applications, LF computational models may use less-accurate turbulence closures, coarser spatial discretizations, looser convergence criteria, or a combination of all three. In contrast, HF models predict the QOI with high accuracy, but at commensurately high cost. Though HF models may be empirical [2,3], the present Paper focuses on computer simulations. The ideal multifidelity scheme combines the strengths of both LF and HF models, generating a faithful surrogate model of the HF QOI at substantially reduced cost. Broadly speaking, LF simulations quickly explore the system's approximate behavior, and a tractable number of (carefully selected) HF runs lends the surrogate model physical accuracy.

Early work in multifidelity modeling was based on the geostatistical method of kriging [4], also known as Gaussian process regression. Kriging, or cokriging in the multivariate case, builds a Gaussian process approximation of the HF model output from multiple, less-expensive observations of lower-fidelity models [5]. This hierarchical approach models the HF output as the sum of a

Presented as Paper 2017-3260 at the 35th AIAA Applied Aerodynamics Conference, Denver, CO, 5–9 June 2017; received 28 April 2018; revision received 22 October 2018; accepted for publication 4 December 2018; published online 31 January 2019. Copyright © 2019 by Ryan W. Skinner. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. All requests for copying and permission to reprint should be submitted to CCC at www.copyright.com; employ the eISSN 1533-385X to initiate your request. See also AIAA Rights and Permissions www.aiaa.org/randp.

*Ph.D. Candidate, Ann and H. J. Smead Department of Aerospace Engineering Sciences. Member AIAA.

†Associate Professor, Ann and H. J. Smead Department of Aerospace Engineering Sciences.

‡Ph.D. Student, Ann and H. J. Smead Department of Aerospace Engineering Sciences.

§Assistant Professor, Ann and H. J. Smead Department of Aerospace Engineering Sciences.

¶Professor, Ann and H. J. Smead Department of Aerospace Engineering Sciences. Fellow AIAA.

**Fernández-Godino, M. G., Park, C., Kim, N. H., and Haftka, R. T., "Review of Multi-Fidelity Models," September 2016, <http://arxiv.org/abs/1609.07196>.

Gaussian process correction term and linearly scaled LF outputs. Methods that are based on Gaussian process regression require a high correlation between the QOI of models with different fidelities, which may not necessarily be the case and which may not be cheap to verify. In addition, rigorous analyses are needed to estimate the required number of HF samples for a successful multifidelity Gaussian process regression, and instead practitioners often stop sampling simply when time or resources are exhausted [6]. Notable applications of cokriging to aerospace applications include building efficient response surfaces with kriging and cokriging [7], incorporation of HF measurement uncertainty into the cokriging framework [8], and using a simplified cokriging covariance matrix to estimate aerodynamic properties of a RAE 2822 airfoil [9], among others.

In recent years, significant progress has been made in extending other well-known UQ techniques to support multiple model fidelities as well. One such extension is the multilevel Monte Carlo (MLMC) method [10,11], which reduces the cost of traditional Monte Carlo (MC) simulation by exploiting successively smaller differences δ between QOI predictions on increasingly finer grids. As δ diminishes under refinement, so does its variance, meaning that finer grids require fewer MC samples. Provided the given black-box solver produces stable solutions to the QOI on coarse grids, MLMC reduces the total cost of computing expectations. Recent work further reduces this cost by using control variates [12,13]. The polynomial chaos expansion (PCE) [14] is another UQ technique that has benefited from extension to a multifidelity setting [15–17]. In particular, the multifidelity PCE approach of Ng et al. [17] constructs a standard PCE of the LF data and subsequently constructs the PCE of an additive correction using a small number of LF and HF samples. The LF PCE and correction PCE are then summed to approximate the HF model.

Though our discussion has so far focused on UQ, multifidelity methods have also proven to be efficient at building surrogate models for design optimization (DO) and optimization under uncertainty (see footnote **) [18]. Examples of multifidelity methods applied to DO include, among others, cokriging via sparse sampling [19], sequential kriging [20], control variates for optimization under uncertainty [21], a PCE-based approach for robust optimization [22], and space mapping [23] applied to transonic airfoil optimization [24].

In addition to multifidelity modeling, and of particular relevance to this Paper, model reduction techniques based on reduced-basis representations have demonstrated utility in reducing total simulation cost without adversely affecting accuracy [25–29]. Recent work provides some insight into why reduced-basis methods are effective, in that the differential equations used to model some physical systems exhibit exploitable low-rank behavior [30]. A low-rank system is characterized by the rapid decay of the largest eigenvalues of the discrete solution. Various UQ schemes leveraging reduced-basis representations of the PDEs have already experienced success in this area [15,31–33].

Our present approach combines a multifidelity framework with a particular reduced-basis representation [32–34], in which each basis corresponds to a unique realization of the system's input parameters. Our technique is relevant to both DO and UQ. Furthermore, it is nonintrusive, meaning it can be applied to any black-box code. Additionally, as long as the modeled system remains low rank in the QOI, the method's cost is formally independent of the dimension of the parameter space. That is, its efficacy does not depend on the number of design variables nor uncertain parameters the system takes as input, provided additional parameters do not substantially raise the system rank. We again emphasize that the modeled system must be low rank in the QOI. As a counterexample, consider a QOI that has parameter-dependent discontinuities, e.g., shocks that change position with Mach number. For such a system, choosing pressure as our QOI is less likely to lead to a low-rank approximation, and therefore the method as presented is unlikely to be effective in this situation.

To ground our discussion, we use the present method to quantify variation in pressure coefficient C_p as a function of geometric parameters defining a two-dimensional (2D) NACA four-digit series airfoil at low α on the C_p distribution across its surface. The input

parameters we vary are the maximum camber m , location of maximum camber p , maximum thickness t , and angle of attack α . Chord length c is fixed at unity. We choose C_p as our QOI because it could be used to calculate many performance metrics, including sectional coefficients of the lift and pitching moment. Our method uses a single HF and a single LF model and can thus be described as bifidelity, although it could be extended to more than two fidelities in a straightforward manner similar to what is done by Narayan et al. [32]. The HF model is a grid-independent Spalart–Allmaras [35] Reynolds-averaged Navier–Stokes (RANS) simulation. To explore the influence of LF model choice, we present results based on two LF models: an Euler flow simulation vs a low-spatial resolution RANS simulation with unresolved boundary layers (BLs).

Once we have picked our QOI and parameter space, the general approach of our bifidelity model is as follows. Assume we determine the $\mathcal{O}(1000)$ realizations that are needed to sufficiently explore the parameter space of our airfoil system. We run one LF simulation for each realization and use the output data to identify both a reduced basis and an interpolation rule mapping the reduced basis back to all $\mathcal{O}(1000)$ LF realizations. Key to our approach is that the bases of reduced-basis approximation each correspond to specific parametric realizations of the airfoil geometry. Second, a select few HF simulations, for example, $\mathcal{O}(10)$, are run for each of these basis-parametric realizations. The bifidelity model then combines the HF results with the LF interpolation rule, generating an approximation of the HF response at all of the other parametric realizations that were only run with the LF model. If the approximation's accuracy is acceptable, the full battery of $\mathcal{O}(1000)$ HF simulations need not be run.

Substantial cost reductions can be realized on the order of $\mathcal{O}(1000/10) \sim 100\times$ in this example, assuming the LF cost is negligible relative to that of the HF model. This cost reduction applies directly to resources spent on design space exploration (e.g., wing geometry and flow control strategy), characterization of system response to environmental inputs (e.g., airspeed and angle of attack), or a combination of both. As long as the LF model sufficiently captures trends in the system's response to design variables, it need not predict system behavior accurately, nor even in an unbiased manner. Our bifidelity modeling approach most benefits systems that have a large parameter space but low-rank behavior of the QOI within this space.

The remainder of the Paper is organized as follows. Section II specifies the models used (one HF and two LF). Section III outlines the solution strategy, detailing the flow simulations, mesh generation, mesh deformation, and cost of LF and HF runs. Section IV presents results, touching on bifidelity model accuracy and reduction in total computational cost for our NACA airfoil application. We conclude in Sec. V with a discussion of results and recommendations for practitioners.

II. Model Specification

A. Flow Simulation and Domain

Though 2D NACA airfoils command little topical interest, this family is nonetheless selected for two reasons. First, it is a standard academic geometry that was thoroughly studied in the 1930s and 1940s, resulting in readily available validation data [36]. Second, applying the bifidelity approach to a NACA airfoil offers a proving ground for the method, which can be scaled up to more industrially relevant aerodynamic systems in future work.

For the present study, we allow $\pm 20\%$ uniform random variation from the NACA 4412 values of maximum camber $m = 0.04$, location of maximum camber $p = 0.4$, and maximum thickness $t = 0.12$. As we shall see, this level of variation leads to a coefficient of variation in C_p greater than 40% at certain airfoil locations. Chord length is fixed at $c = 1$ m. The angle of attack is also sampled randomly and uniformly distributed over the range $\alpha \in [0 \text{ deg}, 6 \text{ deg}]$. Figure 1 helps visualize this geometric variation at zero angle of attack, and Table 1 summarizes our parameter space. Although we employ uniform sampling in this study, finding an optimal sampling strategy to generate the LF snapshots is an




Fig. 1 NACA 4412 airfoil (solid line) with range of surface coordinates for $\pm 20\%$ random variation in parameters m , p , and t . Here, α is fixed at 0° deg, though this study considers $\alpha \in [0^\circ \text{ deg}, 6^\circ \text{ deg}]$ and its subset $\alpha \in [0^\circ \text{ deg}, 2^\circ \text{ deg}]$.

Table 1 Parameter space considered in this study; each parametric realization is drawn from a uniform distribution (note an additional analysis is performed for a subset of the full α domain)

Parameter	Distribution
m	$U[0.032, 0.048]$
p	$U[0.32, 0.48]$
t	$U[0.096, 0.144]$
α	$U[0^\circ \text{ deg}, 6^\circ \text{ deg}]$ and $U[0^\circ \text{ deg}, 2^\circ \text{ deg}]$

interesting open research question, even more so when the inputs do not follow a uniform distribution.

All simulations are incompressible; Euler flow is run in FUN3D [38], and RANS flow is run in PHASTA, a parallel hierarchic adaptive stabilized transient analysis computational fluid dynamics code [39]. For the RANS simulations, we target a Reynolds number of 1.52 million with respect to the airfoil chord. Bifidelity performance will be explored using Euler flow and coarse-mesh RANS as low-fidelity models and mesh-independent RANS as the high-fidelity model. RANS simulations use the Spalart–Allmaras turbulence closure [35].

A schematic of the initial computational fluid domain is shown in Fig. 2. The domain has been fashioned after the work of Diskin et al. [37] and has streamwise length $L = 999$ m and vertical height $H = 998$ m. These dimensions correspond to Cartesian directions x

and y , respectively, and are large to mitigate far-field boundary effects on the near-field solution. Fluid moving in the $+x$ direction enters the domain from the 180° deg, radius 499 m, circular-arc face. A NACA 0012 airfoil of chord length 1.0 m sits in the middle of the domain, which will later be deformed to different NACA geometries. Its leading edge is coincident with the origin. The airfoil has no-slip boundary conditions applied for the RANS simulations and slip-wall boundary conditions for the Euler simulations. Fluid exits the domain on the three outflow faces, one located at $x = 500$ m, and the others on the top and bottom angled at 3° to the x axis.

B. Mesh Generation

For all runs, we begin by meshing the fluid volume described previously, surrounding a standard NACA 0012 airfoil at zero angle of attack. We choose this symmetric geometry instead of a NACA 4412 airfoil in order to replicate the meshing procedure of Diskin et al. [37]. We later deform the symmetric mesh into the NACA 4412-related configurations. A unique mesh is constructed for each model we consider in the bifidelity context: fine-grid RANS (fine RANS), coarse-grid RANS with underresolved boundary layers (coarse RANS), and Euler flow (Euler). Meshes are shown in Fig. 3, and their main characteristics are summarized in Table 2 and detailed in the following. Although Fig. 3 shows 2D renderings of the meshes, 3D versions compatible with our simulation tools are created by extruding the surface mesh in the spanwise direction and then tetrahedronizing the resulting one-element-thick volume mesh. Accordingly, the number of grid points reported in Table 2 is twice what it would be for a true 2D mesh, but the relative costs of HF and LF models remain unchanged. Discussion of typical run times and resource consumption is left to Sec. III.A.

The fine RANS mesh is constructed using the results of Diskin et al. [37] as a reference. The surface mesh size h in the flow direction varies quadratically from $0.0496c$ at the middle of the airfoil to $0.0004c$ at the leading and trailing edges, where c is the chord length. The boundary layer has a stretching ratio of 1.25, and its first point off the wall is 5×10^{-6} m away. This yields $y_{\min}^+ < 1$ in all turbulent regions of the flow. Accordingly, the mesh conforms to Spalart's [40] guidelines. The fine RANS mesh is refined to within 0.1% error for the C_p curves, so as to be an order of magnitude less than the 1–4% validation error of a RANS model [41]. Further refinement is possible, but it would artificially increase the cost reduction afforded by the





Fig. 2 Schematic of computational domain with RANS boundary conditions. Wing will be deformed to NACA 4412 and related configurations after meshing the NACA 0012 configuration following Diskin et al. [37].




Fig. 3 Initial symmetric NACA 0012 meshes used for fine RANS, coarse RANS, and Euler, from left to right. All meshes will be deformed to NACA 4412 and related configurations.

Table 2 Details of low- and high-fidelity models used

Model	Fidelity	Flow equations	Mesh grid points	Cost of one run, s	Cost (% of fine RANS)
Fine RANS	High	S-A RANS	240,861	4980	100
Coarse RANS	Low	S-A RANS	9492	10	0.2008
Euler	Low	Euler	59,361	22	0.4418

bifidelity model, and this level of error is sufficient for engineering purposes. This mesh contains ~ 241 thousand grid points and enables accurate solution of the RANS equations at moderate cost.

The main objective of the two low-fidelity meshes (coarse RANS and Euler) is to obtain a low-cost approximation of the C_p curves. Because our main motivation is cost, less importance is placed on grid refinement. To demonstrate that the bifidelity model is able to provide a useful cost reduction in this particular application, we choose our LF models to have a cost of approximately 0.5% that of the LF model. In practice, the suitability of a given LF model, and therefore the cost reduction it enables, is problem dependent and can be assessed by means such as standard validation error or the a posteriori error estimate derived by Hampton et al. [34] for the resulting bifidelity model. The LF RANS mesh represents a coarsening of the HF RANS mesh until this cost threshold is met. The LF Euler mesh is refined until it no longer satisfies this cost requirement. No claim is made here regarding the quality of the LF meshes since ultimately the adequacy of these meshes is determined by how well they perform in a bifidelity model, discussed in Sec. IV. The following paragraph describes the LF meshes in detail.

The coarse RANS mesh differs from its fine counterpart in three ways. First, the coarse RANS mesh makes no attempt to resolve the shear layer downstream of the trailing edge. Second, instead of fully resolving the boundary layer, the coarse mesh has a first point off the wall at 0.01 m ($y_{\min}^+ \sim 4 \times 10^4$) with the same stretching ratio (1.25). The coarse mesh's first y^+ is already well past the defect layer of a turbulent BL and therefore has no hope of accurately capturing wall effects since it does not apply a wall model. Third, in contrast to the fine RANS surface mesh spacing, the coarse mesh maintains a surface spacing of $0.1c$ except where curvature refinement brings it relatively quickly to $\sim 0.004c$. This mesh does not enable a physically accurate solution, but it contains only $\sim 9\text{ k}$ elements and runs with low cost. For simplicity, the Euler mesh is essentially the same as the fine RANS mesh, but it removes boundary layers and wake refinement features, as they are both viscous phenomena.

Our final step is to deform the NACA 0012 mesh, created according to the guidelines of Diskin et al. [39], into a mesh for each variant of a NACA 4412 airfoil in which we are interested. Recall that a NACA four-digit airfoil is defined by its maximum camber m , location of maximum camber p , and maximum thickness t . These parameters are nondimensionalized by the airfoil chord c , which we take to be unity. The angle of attack α is the final parameter varied by deforming the mesh, rather than by changing the inflow velocity. For each parametric realization, we apply a displacement to nodes on the surface of the NACA 0012 airfoil, represented by the black lines in Fig. 4. A linear elastic structural solver within PHASTA satisfies the Dirichlet displacement constraints and deforms the entire mesh in




Fig. 4 Deforming a NACA 0012 airfoil (red) to an arbitrary NACA four-digit airfoil (blue). Point displacements indicated in black.

accordance with the specified geometric parameters. Positive volumes on BL elements are preserved by enforcing an inverse relationship between element volume and stiffness, through a Jacobian-based stiffening approach [42].

III. Solution Strategy

A. Solution Procedure and Efficiency

Numerical solution of the flow equations proceeds until steady results are obtained. A converged NACA 0012 flow solution at $\alpha = 0$ deg acts as the starting point in PHASTA for each new combination of geometric parameters to be tested. After deforming the mesh for each NACA 4412-based geometric realization, we seek a steady solution to the quantity of interest by running the simulation at a large time step of $\Delta t = 12$ chord flights with first-order time integration and a single nonlinear iteration per time step. One chord flight is the transit time defined as the chord length c divided by the freestream speed. Runs continue until the maximum change in C_p over the last 240 chord flights is less than 0.01%, which is another order of magnitude smaller than the 0.1% uncertainty in C_p due to mesh convergence discussed in Sec. II.B. This typically takes around 100 time steps but depends weakly on the mesh and physical model.

For Euler runs, the deformed geometry is still calculated in PHASTA, but then the geometry is passed as input to FUN3D. FUN3D is run with a stopping tolerance of 1×10^{-8} (i.e., conclude the simulation when the root mean square of the residual of every solved equation drops below this value). This criterion is typically met after 250–350 time steps.

When an individual run meets the stopping criterion, C_p information is extracted from the solution file through a ParaView script and stored with an identifier linking the run to its specific realization of geometric parameters. This step can be made part of the solver by marking nodes in a preprocessing step for computation within PHASTA or by directly linking to the ParaView Catalyst library, which has been shown to give flexible, scalable, and efficient extraction of almost any quantity of interest [43,44].

The exact points at which C_p is measured are defined relative to the chord line, meaning data are extracted in a frame in which the airfoil is rotated back to zero α . Whereas measurements' x coordinates are constant across all three meshes and are independent of geometric variation in the airfoil, their y coordinates depend on both the mesh and airfoil shape. There is a combined total of $n = 100$ sampling locations (x_j, y_j) on the upper and lower surfaces. The values of x_j are spaced quadratically in x [i.e., with positions relative to the chord length of $x_j/c = (2j/n)^2$, $j = 0, 1, \dots, n/2 + 1$ for even n]. Each y_j is defined by where the upper or lower surface intersects $x = x_j$. We choose this distribution to place more measurements in areas where gradients of C_p are large, thereby capturing C_p behavior faithfully along both the high-curvature leading edge and the main lifting surfaces of the airfoil. However, the choice of measurement locations is problem dependent and is ultimately left to the practitioner.

All components that make up a run have now been discussed; a brief conversation about computational cost is pertinent. As shown in Table 2, a single Euler run takes 22 s, a coarse RANS run takes 10 s, and a fine RANS run takes 4980 s (roughly two orders of magnitude longer). All figures are CPU time quoted running on an Intel Xeon E5-2650 v3 (Haswell) CPU at 2.30 GHz. Jobs are run simultaneously as they have no interdependence, greatly compressing the time to obtain results needed for the bifidelity model.

B. Bifidelity Model

We construct the bifidelity model by first sampling our parameter space with a chosen LF model and then supplementing our approximate knowledge of system behavior with carefully selected HF runs. The mathematics leverage the rank-revealing QR algorithm [45,46], following the approach of Doostan et al. [47] and Hampton et al. [34] In addition to describing the mathematics in this section, Sec. IV includes figures that illustrate specific steps of the bifidelity modeling process.

The bifidelity method uses the solution to a cheaper, lower-fidelity model of the system of interest to identify a reduced basis and an

interpolation rule for approximating the expensive, high-fidelity solution as a function of its inputs. Formally defining the parameter space, consider the set of random geometric and angle of attack parameters, represented as the random vector

$$\mathbf{y} = [m \quad p \quad t \quad \alpha]^T \quad (1)$$

Each realization of random parameters is denoted by $\mathbf{y}^{(i)}$, where $i = 1, \dots, K$. The values of C_p at $n = 100$ locations on the airfoil surface are collected in a solution vector $\mathbf{u}(\mathbf{y}^{(i)})$ for each i th parametric realization. LF and HF solutions are denoted by a subscript, respectively, $\mathbf{u}_L(\mathbf{y}^{(i)})$ and $\mathbf{u}_H(\mathbf{y}^{(i)})$.

The set of K LF solutions is collected into a matrix:

$$\mathbf{L} \equiv \begin{bmatrix} | & | & | & | \\ \mathbf{u}_L(\mathbf{y}^{(1)}) & \mathbf{u}_L(\mathbf{y}^{(2)}) & \dots & \mathbf{u}_L(\mathbf{y}^{(K)}) \\ | & | & | & | \end{bmatrix}_{n \times K} \quad (2)$$

Each column holds a single LF solution vector, e.g., pressure on the airfoil. A practitioner could accommodate two or more QOI by simply stacking the fields to create larger realization vectors $\mathbf{u}$, with the understanding that this additional complexity could raise the rank of the bifidelity approximation. A rank- r interpolating decomposition (ID) $\hat{\mathbf{L}}$ is constructed, approximating $\mathbf{L}$ as the product of the so-called column skeleton matrix $\mathbf{L}^C$ and the coefficient matrix $\mathbf{\Lambda}$ [47],

$$\begin{aligned} \mathbf{L} \approx \hat{\mathbf{L}} &= \begin{bmatrix} | & | & | & | \\ \mathbf{u}_L(\mathbf{y}^{(i_1)}) & \mathbf{u}_L(\mathbf{y}^{(i_2)}) & \dots & \mathbf{u}_L(\mathbf{y}^{(i_r)}) \\ | & | & | & | \end{bmatrix}_{n \times r} [\Lambda_{ij}]_{r \times K} \\ &= \mathbf{L}^C \mathbf{\Lambda} \end{aligned} \quad (3)$$

It is important to note that the $r \ll K$ basis vectors $\{\mathbf{u}_L(\mathbf{y}^{(i_\ell)})\}_{\ell=1}^r$ are some realizations of the LF solution and form a reduced basis for the approximation of the rest of the LF realizations. The basis configurations are indexed by (i_ℓ) instead of (ℓ) because the vectors chosen for the basis are not necessarily the first r columns of $\mathbf{L}$, which would be enumerated (1) through (r) . The ID identifies the reduced basis using the rank revealing QR algorithm, which is explained in greater mathematical detail in the Appendix. Furthermore, the decomposition rank r is not fixed; it can be varied depending on the desired accuracy of the ID. For a given accuracy, systems that are low-rank in the QOI will permit the use of a smaller reduced basis and thus a smaller r .

To compute the bifidelity model, we next obtain high-fidelity solutions corresponding to the r geometric realizations $\{\mathbf{y}^{(i_\ell)}\}_{\ell=1}^r$ deemed important by the LF ID. Again, the QOI from these r high-fidelity simulations populate a HF column skeleton matrix $\mathbf{H}^C$. If the system is low rank in the QOI, r will be small, and only a handful of HF simulations will be needed to construct $\mathbf{H}^C$. We then reuse the ID to calculate an approximation $\hat{\mathbf{H}}$ of K HF solutions,

$$\begin{aligned} \hat{\mathbf{H}} &= \begin{bmatrix} | & | & | & | \\ \mathbf{u}_H(\mathbf{y}^{(i_1)}) & \mathbf{u}_H(\mathbf{y}^{(i_2)}) & \dots & \mathbf{u}_H(\mathbf{y}^{(i_r)}) \\ | & | & | & | \end{bmatrix}_{n \times r} [\Lambda_{ij}]_{r \times K} \\ &= \mathbf{H}^C \mathbf{\Lambda} \end{aligned} \quad (4)$$

In summary, the bifidelity model requires K low-fidelity runs of which the column snapshots of the QOI are used to identify the r

high-fidelity runs and the coefficient matrix $\mathbf{\Lambda}$ that weights them to provide the bifidelity prediction of the QOI for all parameter combinations of interest. The LF model's low cost allows us to obtain $K \gg r$ samples. We also require $K \gg r$ samples so the ID can recover the best r bases from the LF model.

Normally, $\mathbf{H}$ is unavailable because it requires expensive HF solutions for all parametric realizations. This high cost is why in practice we only require r HF model evaluations to construct the bifidelity model. However, for this problem, we perform not only the required r but all K HF simulations. While this would not be done in applications of interest, it is tractable for this NACA airfoil problem and allows for exact evaluation of bifidelity model performance. To this end, we will report the relative error of the bifidelity model

$$\epsilon \equiv \|\hat{\mathbf{H}} - \mathbf{H}\| / \|\mathbf{H}\| \quad (5)$$

in the matrix 2-norm after excluding basis configurations. Since $\hat{\mathbf{H}}$ is constructed through an interpolating decomposition, the r columns corresponding to basis configurations are reproduced without error. Including those columns would artificially reduce the relative error measure; this effect would grow with increasing r . Because ϵ is computed with all nonbasis configurations, it can be thought of as a validation error with the largest sample size possible. As is commonly done in the statistical regression literature, estimates of ϵ may be computed with a smaller number of extra HF model evaluations [48]. Finally, we also define the true relative error in the matrix 2-norm between LF and HF predictions of the QOI as

$$\epsilon_{\text{LVH}} \equiv \|\mathbf{L} - \mathbf{H}\| / \|\mathbf{H}\| \quad (6)$$

IV. Results and Discussion

In this Paper, we simultaneously compare two BF models, and thus we require two sets of low-fidelity runs. $K = 500$ runs are conducted for both Euler (first LF) and coarse RANS (second LF) models. The ID and corresponding BF model are computed for $r = 1, 2, \dots, 80$ for each choice of LF model. In practice, we only require r HF model evaluations to construct each of the two BF models. Though not required for the two BF models, to allow the exact evaluation of BF model performance, we perform not r but all K HF simulations. As previously described in Sec. III.B, while this would not be done in applications of interest, it is tractable for this NACA airfoil problem.

The cost reduction afforded by the BF model depends directly on the rank of the QOI, with lower-rank behavior permitting greater cost reduction. To assess whether our system is low rank in its QOI, singular values [49] of C_p for our models are plotted in Fig. 5. Rapid singular value decay suggests that fewer basis configurations, and thus fewer costly HF model evaluations, will be needed to construct the BF model. We emphasize that, in practice, we do not have access to the entire $\mathbf{H}$ matrix so we do not know its singular values. We do, however, have the entire $\mathbf{L}$ matrix, which must be low rank if the BF model is to offer substantial cost reduction over a full battery of HF simulations. Here, we show singular values of both $\mathbf{L}$ (one for each LF model) and $\mathbf{H}$ as a means of illustrating the rank of our chosen models. For all models,




Fig. 5 Normalized singular values of the data matrices $\mathbf{L}$ (one for each LF model) and $\mathbf{H}$. Two sets of singular values are given for each, corresponding to the full and restricted α domains.

we see that the normalized singular value magnitude drops by three orders of magnitude by index 10 and then decays more slowly, dropping another three orders of magnitude by index 70. This behavior is not truly low rank, but we will see later that it is low-rank enough to produce BF models with a useful ratio of cost reduction to accuracy.

Concluding our discussion of Fig. 5, two sets of singular values are shown for each model: one corresponds to all 500 realizations with $\alpha \in [0 \text{ deg}, 6 \text{ deg}]$, and one corresponds to the subset of 168 runs satisfying $\alpha \in [0 \text{ deg}, 2 \text{ deg}]$. These will be referred to as the full and restricted α domains, respectively. For the remainder of this Paper, we consider both as a means of varying the solution rank. However as can be seen, the change in singular value magnitude between α domains is relatively small, meaning some high-rank behavior persists due to the other input parameters defining airfoil shape.

We now turn to the visualization of the BF modeling process shown in Fig. 6. This example uses the coarse RANS LF model with the full α domain. We choose $r = 10$ for clarity of presentation and emphasize that a rank-10 BF model requires only ten realizations of the HF model. Subfigure a corresponds to Eq. (3), showing the full set of aggregate LF data $\mathbf{L}$ and highlighting the C_p curves associated with the basis vectors in $\mathbf{L}^C$ identified by the ID. Subfigure b demonstrates rerunning the LF basis configurations with the HF model and compares the C_p curves in $\mathbf{L}^C$ and $\mathbf{H}^C$. Here, we draw attention to the large discrepancies between LF and HF model results. Subfigure c corresponds to Eq. (4), wherein the BF model $\hat{\mathbf{H}}$ is constructed from the HF column skeleton $\mathbf{H}^C$ and the LF coefficient matrix $\mathbf{A}$. Finally, subfigure d evaluates the BF model accuracy by plotting standard C_p curves obtained from the LF, HF, and BF models for three representative parametric realizations.

Note that data in Figs. 6a–6c are plotted against the normalized airfoil surface location, which shows variability in aggregate data near the leading edge more clearly than a standard mapping of x/c to C_p . This location on airfoil is calculated as the normalized streamwise distance from the airfoil's trailing edge. Moving clockwise around the airfoil is interpreted as the positive direction, so positive values of location on airfoil correspond to the pressure surface, while negative values correspond to the suction surface.

We also visualize the BF modeling process for the remaining combinations of LF model and α domain. Figure 7 corresponds to the

Euler LF model and full α domain. Figure 8 corresponds to the coarse RANS LF model and restricted α domain. Lastly, Fig. 9 corresponds to the Euler LF model and restricted α domain.

There are two main conclusions to be drawn from Figs. 6–9. First, subfigures b show the large discrepancy between LF and HF predictions of C_p . Because we have access to $\mathbf{H}$ (due to the low cost of the HF system), as well as the LF model-specific $\mathbf{L}$ matrices, we can compute the relative error ϵ_{LVH} between all LF and HF predictions using Eq. (6). For the coarse RANS LF model, the full and restricted α domain cases have $\epsilon_{\text{LVH}} = 41.9\%$ and $\epsilon_{\text{LVH}} = 34.4\%$, respectively. For the Euler LF model, these values drop to $\epsilon_{\text{LVH}} = 18.2\%$ and $\epsilon_{\text{LVH}} = 16.7\%$, respectively. Furthermore, the two LF models predict very different C_p curves from one another, which can be seen by comparing Fig. 6 to Fig. 7 and also Fig. 8 to Fig. 9. Despite these differences in LF model predictions, subfigures d in Figs. 6–9 demonstrate that their BF results are visually indistinguishable from HF results, at least for the three representative parametric realizations shown. This is the first indication that as long as the LF model captures the correct trends in the QOI under parametric variation, the absolute accuracy of the LF QOI prediction at any single parametric realization is inconsequential to the BF model's accuracy.

To formalize this assessment of BF model accuracy, we now consider the true relative error metric ϵ defined in Eq. (5). Figure 10 plots ϵ as a function of model rank r for BF models constructed from different LF models and for different choices of the α domain. Figure 10 also includes the relative error (defined analogously to ϵ_{LVH}) of the best rank- r approximation of $\mathbf{H}$ computed via the singular value decomposition (SVD) algorithm [49] and independently of the LF model. Henceforth, we will refer to this error as the optimal relative error. It is a benchmark for evaluating BF accuracy, but it is not to be confused with the best possible accuracy of the BF model. There is no guarantee, nor even expectation, that any of the r bases identified by the best rank- r approximation to $\mathbf{H}$ will be physical, in that their C_p curves correspond to some NACA airfoil within our parameter space. Optimal relative errors in Fig. 10 that continue to drop by orders of magnitude with increasing rank imply that our NACA airfoil system is receptive to low-rank approximation, insofar as the prediction of C_p is




Fig. 6 Bifidelity model with rank $r = 10$ and $\alpha \in [0 \text{ deg}, 6 \text{ deg}]$ for coarse RANS LF and fine RANS HF models. Relative error of the bifidelity model is 1.97%.




Fig. 7 Bifidelity model with rank $r = 10$ and $\alpha \in [0 \text{ deg}, 6 \text{ deg}]$ for Euler LF and Fine RANS HF models. Relative error of the bifidelity model is 1.56%.




Fig. 8 Bifidelity model with rank $r = 10$ and $\alpha \in [0 \text{ deg}, 2 \text{ deg}]$ for coarse RANS LF and fine RANS HF models. Relative error of the bifidelity model is 0.94%.

concerned. This is consistent with our previous discussion of the models' singular value structure (Fig. 5).

The BF relative errors ϵ reported in Fig. 10 constitute a marked improvement over the 15–40% errors ϵ_{LVH} of each LF model alone. Note that ϵ drops below 1% for ranks between $r = 7$ and $r = 21$, with the exact behavior depending on LF model and α domain. For large r , the BF model accuracy plateaus at between 0.1 and 0.3% error. This error cannot be improved beyond a certain level due to discrepancies between the LF and HF models. For small $r < 10$, our method identifies a

sufficiently good parametric basis $\{\mathbf{y}^{(i)}\}_{i=1}^r$ and coefficient matrix $\mathbf{A}$ such that the BF model's relative error is within an order of magnitude of the optimal relative error. As r increases, however, the BF model is not able to maintain this gap due to discrepancies between LF and HF models. In short, the LF model captures major trends in the HF model well, but as we include more nuanced behavior by increasing r , the inaccuracies of our LF model cause the BF model accuracy to plateau.

An equally important aspect of the BF model is its potential for cost savings. To rigorously assess this, we define cost savings as the cost




Fig. 9 Bifidelity model with rank $r = 10$ and $\alpha \in [0 \text{ deg}, 2 \text{ deg}]$ for Euler LF and fine RANS HF models. Relative error of the bifidelity model is 0.55%.




Fig. 10 Relative error vs rank for the two bifidelity models and the best rank- r approximations of the HF (fine RANS) data.

of a full battery of K HF runs (each one being C_{HF}) divided by the cost C_{BF} of the BF model. The latter is equal to the total cost of K LF runs (each one being C_{LF}) and r HF runs of the basis configurations, giving

$$\text{cost savings} = \frac{KC_{HF}}{C_{BF}} = \frac{KC_{HF}}{KC_{LF} + rC_{HF}} \quad (7)$$

where C_{LF} and C_{HF} are the costs of individual LF and HF runs found in Table 2. We report cost savings in units of $\times$ (read “times”). That is, a cost savings of $10\times$ means that the computational cost to construct $\mathbf{H}$ with a HF model by brute force is 10 times greater than the cost to construct $\hat{\mathbf{H}}$ using the bifidelity framework. Table 3 shows cost savings as a function of BF model rank. The two sets of columns

Table 3 Cost savings, Eq. (7), based on the full ($K = 500$ runs) and restricted ($K = 168$ runs) α domains

Rank	$\alpha \in [0 \text{ deg}, 6 \text{ deg}]$		$\alpha \in [0 \text{ deg}, 2 \text{ deg}]$	
	Euler	Coarse RANS	Euler	Coarse RANS
10	49.78 $\times$	49.90 $\times$	16.73 $\times$	16.77 $\times$
20	24.89 $\times$	24.94 $\times$	8.36 $\times$	8.38 $\times$
30	16.59 $\times$	16.63 $\times$	5.58 $\times$	5.59 $\times$
LF model only	224.3 $\times$	498 $\times$	224.3 $\times$	498 $\times$

correspond to the full and restricted α domains and thus have different values of $K = 500$ and $K = 168$, respectively. There is little difference between the Euler and coarse RANS columns for a given α domain because for both LF models $C_{LF} \ll C_{HF}$. Finally, the row labeled *LF Model Only* corresponds to $r = 0$, where a BF model is not constructed at all.

BF model performance is a combination of accuracy and cost savings. Both must be sufficiently high for the BF model to be useful, but while the former increases with r , the latter decreases. We highlight this inherent performance tradeoff with three examples taken from Fig. 10 and Table 3. First, the rank-30 BF model for the restricted α domain constructed from the Euler LF model attains 0.21% error at a cost savings of $6\times$. Second, the rank-20 BF model for the full α domain constructed from the Euler LF model attains 0.58% error at a cost savings of $25\times$. Third, the rank-10 BF model for the full α domain constructed from the coarse RANS LF model attains 2.0% error at a cost savings of $50\times$. Despite these performance tradeoffs, our BF model predicts the pressure coefficient C_p at a cost that is often more than one order of magnitude cheaper than a full suite of HF simulations.

From Fig. 10, we conclude that each of the BF models provides a solution to the full design space that has less than 1% error while expending resources comparable to running only a small subset of 7–21 HF runs. This reduces the total simulation cost vs a full suite of 168–500 HF runs by about $20\times$. Practically, these accuracies are well

within the 1–4% validation error of RANS models applied to the comparable case of a 2D NACA 0012 airfoil at $\alpha = 0^\circ$ and $Re = 6$ Million [41]. Thus, our bifidelity model is more than adequate for design space exploration of this physical system to within modeling error.

In UQ, practitioners seek to determine how uncertainty in system inputs influences system outputs. The nonintrusive approach to this question treats the simulation as a black box, first sampling the output for many realizations of the uncertain input parameters and then computing statistics to describe the output distribution. The main advantage of the nonintrusive approach is not having to modify existing simulation code. Our bifidelity model reduces the cost of the first step in nonintrusive UQ: mapping sampled input parameters to system outputs. Though the parameter range used in this study is much larger than the uncertainties a practitioner would usually encounter, Fig. 11 nonetheless highlights our bifidelity model's applicability to UQ by computing statistics of the C_p distribution computed from the LF, HF, and BF models at each location on the airfoil where data are sampled. As we have seen in Figs. 5 and 10, a smaller parameter space may reduce the rank of the system, meaning that in a true UQ application and for a fixed total computational cost, the accuracy of the BF model may be much higher.

In Fig. 11, the mean of C_p is plotted alongside regions representing twice the standard deviation $\pm 2\sigma$ from the mean. Subfigures a–b and c–d correspond to the full and restricted α domains, respectively. Subfigures a and c and b and d correspond to the coarse RANS and Euler LF models. Furthermore, the combinations of LF model and α domain in subfigures a–d are the same as those in Figs. 6–9. In each subfigure, statistics are shown for the LF model's $\mathbf{L}$, the HF model's $\mathbf{H}$, and the rank-10 BF model's $\hat{\mathbf{H}}$ based on that LF model. In each case, the difference between the statistics of the LF and HF models is substantial, whereas the BF and HF models' mean and $\pm 2\sigma$ regions are visually indistinguishable. Also note that in subfigures a and b the coefficient of variation in C_p is greater than 40% at certain airfoil locations. Because the BF model accurately predicts individual pressure curves, its excellent prediction of C_p statistics is not particularly surprising due to statistical moments being integral quantities. Thus, the utility of our bifidelity model extends to accurately predicting statistics of the QOI at substantially reduced cost.

V. Conclusions

This Paper applies a reduced-basis multifidelity model to explore the parameter space of a 2-D NACA 4412 airfoil subject to variation in geometric parameters and angle of attack. The goal is to avoid a full battery of high-fidelity simulations, which is not practical at scale for high-complexity aerodynamic systems. The approach is to use a low-fidelity model to generate an approximation of the high-fidelity (HF) solutions when supplemented by a small number of carefully chosen HF runs. The low-fidelity (LF) samples are used to identify a reduced basis along with an interpolation rule to generate an approximation of HF solutions corresponding to arbitrary realizations of the airfoil geometry. Because one LF and one HF model are used to construct the approximation of HF solutions, our multifidelity model is technically a bifidelity model, although it would be straightforward to extend to more than two fidelities [32].

For the BF model to be successful, it must predict the HF results with high accuracy while substantially reducing the total computational cost. Success is predicated upon the physical system being low rank in the quantity of interest (QOI). Furthermore, the LF system must respond in a qualitatively similar way to the HF system, at least within the parameter space considered for the QOI. The LF models used in this study are Euler flow and coarse-mesh Reynolds-averaged Navier–Stokes (RANS), whereas the HF model is fine-mesh RANS that predicts C_p curves within 0.1% error. This gives one order of magnitude separation from the validation error of RANS models applied to the comparable case of a 2D NACA 0012 airfoil at $\alpha = 0^\circ$ and $Re = 6$ Million [41].

We also consider two angle of attack ranges as a means of varying the size of the parameter space. In both cases, the LF system is low rank enough for the BF model to be successful. By low rank, we mean that the LF system's largest singular values decay rapidly as in Fig. 5. Accordingly, we expect any reasonable approximation in a reduced basis, including the employed BF model, to be accurate. Because the restricted α -domain singular values decay slightly more rapidly than their full α -domain counterparts, we expect their low-rank approximations to perform slightly better. Lack of large differences in the singular values between α domains is attributed to significant complexity remaining in the other NACA airfoil parameters.




Fig. 11 Statistics of C_p computed for the LF, HF, and rank-10 BF models, for different combinations of the α domain and LF model.

Our numerical examples illustrate that the BF model predicts C_p along the airfoil with less than 1% error at a cost that is roughly 20× less than a full suite of HF simulations. This is well within the 1–4% RANS validation error [41] mentioned previously. Practically, our BF model is thus more than adequate for design space exploration of this NACA airfoil system to within modeling error. The BF model also predicts statistics of C_p (mean and standard deviation) at each point along the airfoil that are visually indistinguishable from the HF results. Though not particularly surprising due to its high accuracy when predicting individual configurations, we show that our model is equally applicable to uncertainty quantification by predicting statistics of the QOI at substantially reduced cost.

In practice, there are two ways that we envision assessing the quality of the BF approximation. A standard approach would be to employ widely used statistical tools such as validation error computed based on additional HF samples of nonbasis configurations. This is commonly done in the statistical regression literature [48]. In practice, this means picking a handful of configurations either at random or based on maximal parametric distance in the Euclidean norm to the basis configurations, running them through the HF model, and computing the relative error of the BF model from this small sample of additional runs. If the practitioner seeks an optimal design, it would be natural to check points near the optimal configuration predicted by the bifidelity model. Additionally, the recent work of Hampton et al. [34] develops an a posteriori bound on the error of the BF approximation presented here.

Ultimately, three major factors influence bifidelity model performance. First, the QOI must be low rank within the parameter space. Even within the same physical system, different QOI may have different ranks, and different LF models may be more appropriate for predicting different QOI. Second, rank depends on the parameter space. Generally, the larger the parameter space, the more complicated the behavior of the QOI will be, and the higher its rank will climb. Third, the cost savings depends on the fraction of configurations for which the QOI is estimated by the bifidelity model rather than directly run through the HF model. In terms of Eq. (7), the larger the K and the smaller the r , the higher the cost savings. Such an analysis is strongly problem dependent, hinging on both the number of realizations required to sufficiently explore the parameter space and the rank of the QOI on that space. The bifidelity approach is thus most useful when the parameter space is large enough to require large K , but the response is low rank, and a LF model is still able to capture relevant trends.

Appendix: The Rank-Revealing QR Algorithm

The ID, $L \approx \hat{L} = L^C \Lambda$, is constructed using the rank-revealing QR algorithm. In particular, the rank- r approximation $\hat{L}$ to $L \in \mathbb{R}^{n \times K}$ is obtained using column-pivoted (truncated) QR factorization,

$$\begin{aligned} LP &\approx Q[R_{11} \quad R_{12}] \\ &= QR_{11}[I \quad R_{11}^\dagger R_{12}] \end{aligned} \quad (A1)$$

where $Q \in \mathbb{R}^{n \times r}$ has r orthogonal columns, $R_{11} \in \mathbb{R}^{r \times r}$ is upper triangular, $R_{12} \in \mathbb{R}^{r \times K-r}$, $P \in \mathbb{R}^{K \times K}$ is a permutation matrix, and † denotes the (e.g., Moore–Penrose) pseudoinverse. It is straightforward to show that QR_{11} is equal to the column skeleton L^C , and the indices of columns from L that actually comprise L^C can be identified from the permutation matrix P . Finally, the coefficient matrix Λ is equal to $[I \quad R_{11}^\dagger R_{12}]P^T$. The inquisitive reader is directed to the works of Gu and Eisenstat [45] and Cheng et al. [46] for a thorough discussion of how the rank-revealing QR algorithm computes this factorization.

Acknowledgments

Ryan W. Skinner's research for this Paper was conducted with U.S. Government support under contract FA9550-11-C-0028 and awarded by the Department of Defense, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate Fellowship, 32 CFR 168a. This material is based upon work of Alireza Doostan supported by the U.S. Department of Energy Office

of Science, Office of Advanced Scientific Computing Research, under award number DE-SC0006402 and NSF grant CMMI-1454601. Eric L. Peters and John A. Evans were partially supported by the Air Force Office of Scientific Research under grant number FA9550-14-1-0113. Kenneth E. Jansen and John A. Evans were partially supported by National Science Foundation grant CBET-1710670. An award of computer time was provided by the Innovative and Novel Computational Impact on Theory and Experiment program. This research used resources of the Argonne Leadership Computing Facility, which is a U.S. Department of Energy Office of Science User Facility supported under contract DE-AC02-06CH11357. The authors gratefully acknowledge Simmetrix, Inc., for its meshing and geometric modeling libraries, Altair Engineering for its linear algebra solver library, and Kitware for its visualization tools. Lastly, the authors thank the Scientific Computation Research Center (SCOREC) at Rensselaer Polytechnic Institute (RPI) for its core mesh partitioning tools, supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under award DE-SC00066117 (FASTMath SciDAC Institute).

References

- [1] Peherstorfer, B., Willcox, K., and Gunzburger, M., "Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization," 2016, <https://epubs.siam.org/doi/10.1137/16M1082469>.
- [2] Forrester, A. I. J., "Black-Box Calibration for Complex-System Simulation," *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Vol. 368, No. 1924, 1924, Aug. 2010, pp. 3567–3579. doi:10.1098/rsta.2010.0051
- [3] Kuya, Y., Takeda, K., Zhang, X., and Forrester, A. I. J., "Multifidelity Surrogate Modeling of Experimental and Computational Aerodynamic Data Sets," *AIAA Journal*, Vol. 49, No. 2, 2011, pp. 289–298. doi:10.2514/1.J050384
- [4] Krige, D. G., "A Statistical Approach to Some Basic Mine Valuation Problems on the Witwatersrand," *Journal of the Southern African Institute of Mining and Metallurgy*, Vol. 52, No. 9, 1952, pp. 201–203.
- [5] Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P., "Design and Analysis of Computer Experiments," *Statistical Science*, Vol. 4, No. 4, Nov. 1989, pp. 409–423. doi:10.1214/ss/1177012413
- [6] Forrester, A. I. J., and Keane, A. J., "Recent Advances in Surrogate-Based Optimization," *Progress in Aerospace Sciences*, Vol. 45, No. 1, Jan. 2009, pp. 50–79. doi:10.1016/j.paerosci.2008.11.001
- [7] Laurenceau, J., and Sagaut, P., "Building Efficient Response Surfaces of Aerodynamic Functions with Kriging and Cokriging," *AIAA Journal*, Vol. 46, No. 2, 2008, pp. 498–507. doi:10.2514/1.32308
- [8] Qian, P. Z. G., and Jeff Wu, C. F., "Bayesian Hierarchical Modeling for Integrating Low-Accuracy and High-Accuracy Experiments," *Technometrics*, Vol. 50, No. 2, 2008, pp. 192–204. doi:10.1198/004017008000000082
- [9] Han, Z., Zimmerman, R., and Görtz, S., "Alternative Cokriging Method for Variable-Fidelity Surrogate Modeling," *AIAA Journal*, Vol. 50, No. 5, 2012, pp. 1205–1210. doi:10.2514/1.J051243
- [10] Heinrich, S., "Multilevel Monte Carlo Methods," *In Large-Scale Scientific Computing*, Lecture Notes in Computer Science, Springer–Verlag, Berlin, June 2001, pp. 58–67. doi:10.1007/3-540-45346-6_5
- [11] Giles, M. B., "Multilevel Monte Carlo Path Simulation," *Operations Research*, Vol. 56, No. 3, June 2008, pp. 607–617. doi:10.1287/opre.1070.0496
- [12] Nobile, F., and Tesei, F., "A Multi-Level Monte Carlo Method with Control Variate for Elliptic PDEs with Log-Normal Coefficients," *Stochastic Partial Differential Equations: Analysis and Computations*, Vol. 3, No. 3, Sept. 2015, pp. 398–444. doi:10.1007/s40072-015-0055-9
- [13] Fairbanks, H. R., Doostan, A., Ketelsen, C., and Iaccarino, G., "A Low-Rank Control Variate for Multilevel Monte Carlo Simulation of High-Dimensional Uncertain Systems," *Journal of Computational Physics*, Vol. 341, July 2017, pp. 121–139, <https://www.sciencedirect.com/science/article/pii/S0021999117302619>.

- [14] Najm, H. N., "Uncertainty Quantification and Polynomial Chaos Techniques in Computational Fluid Dynamics," *Annual Review of Fluid Mechanics*, Vol. 41, Jan. 2009, pp. 35–52.
doi:10.1146/annurev.fluid.010908.165248
- [15] Doostan, A., Ghanem, R. G., and Red-Horse, J., "Stochastic Model Reduction for Chaos Representations," *Computer Methods in Applied Mechanics and Engineering*, Vol. 196, No. 37, Aug. 2007, pp. 3951–3966.
doi:10.1016/j.cma.2006.10.047
- [16] Ghanem, R., Saad, G., and Doostan, A., "Efficient Solution of Stochastic Systems: Application to the Embankment Dam Problem," *Structural Safety*, Vol. 29, No. 3, July 2007, pp. 238–251.
doi:10.1016/j.strusafe.2006.07.015
- [17] Ng, L. W. T., and Eldred, M., "Multifidelity Uncertainty Quantification Using Non-Intrusive Polynomial Chaos and Stochastic Collocation," *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, Structural Dynamics, and Materials and Co-located Conferences*, AIAA Paper 2012-1852, April 2012.
doi:10.2514/6.2012-1852
- [18] Queipo, N. V., Haftka, R. T., Shyy, W., Goel, T., Vaidyanathan, R., and Tucker, P. K., "Surrogate-Based Analysis and Optimization," *Progress in Aerospace Sciences*, Vol. 41, No. 1, Jan. 2005, pp. 1–28.
doi:10.1016/j.paerosci.2005.02.001
- [19] Forrester, A. I. J., Söbester, A., and Keane, A. J., "Multi-Fidelity Optimization Via Surrogate Modelling," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, Vol. 463, No. 2088, Dec. 2007, pp. 3251–3269.
doi:10.1098/rspa.2007.1900
- [20] Huang, D., Allen, T. T., Notz, W. I., and Miller, R. A., "Sequential Kriging Optimization Using Multiple-Fidelity Evaluations," *Structural and Multidisciplinary Optimization*, Vol. 32, No. 5, Nov. 2006, pp. 369–382.
doi:10.1007/s00158-005-0587-0
- [21] Ng, L. W. T., and Willcox, K. E., "Multifidelity Approaches for Optimization Under Uncertainty," *International Journal for Numerical Methods in Engineering*, Vol. 100, No. 10, Dec. 2014, pp. 746–772.
doi:10.1002/nme.v100.10
- [22] Padron, A. S., Alonso, J. J., and Eldred, M. S., "Multi-Fidelity Methods in Aerodynamic Robust Optimization," *18th AIAA Non-Deterministic Approaches Conference*, AIAA Paper 2016-0680, 2016.
doi:10.2514/6.2016-0680
- [23] Bandler, J. W., Cheng, Q. S., Dakrouy, S. A., Mohamed, A. S., Bakr, M. H., Madsen, K., and Sondergaard, J., "Space Mapping: The State of the Art," *IEEE Transactions on Microwave Theory and Techniques*, Vol. 52, No. 1, Jan. 2004, pp. 337–361.
doi:10.1109/TMTT.2003.820904
- [24] Leifsson, L., and Koziel, S., "Multi-Fidelity Design Optimization of Transonic Airfoils Using Physics-Based Surrogate Modeling and Shape-Preserving Response Prediction," *Journal of Computation Science*, Vol. 1, No. 2, June 2010, pp. 98–106.
doi:10.1016/j.jocs.2010.03.007
- [25] Amsallem, D., Zahr, M. J., and Farhat, C., "Nonlinear Model Order Reduction Based on Local Reduced-Order Bases," *International Journal for Numerical Methods in Engineering*, Vol. 92, No. 10, Dec. 2012, pp. 891–916.
doi:10.1002/nme.v92.10
- [26] Rozza, G., Huynh, D. B. P., and Patera, A. T., "Reduced Basis Approximation and a Posteriori Error Estimation for Affinely Parametrized Elliptic Coercive Partial Differential Equations," *Archives of Computational Methods in Engineering*, Vol. 15, No. 3, Sept. 2008, p. 229.
doi:10.1007/s11831-008-9019-9
- [27] Maday, Y., and Rønquist, E. M., "A Reduced-Basis Element Method," *Journal of Scientific Computing*, Vol. 17, Nos. 1–4, Dec. 2002, pp. 447–459.
doi:10.1023/A:1015197908587
- [28] Quarteroni, A., Manzoni, A., and Negri, F., *Reduced Basis Methods for Partial Differential Equations: An Introduction*, Springer–Verlag, Cham, Aug. 2015.
- [29] Jofre, L., Geraci, G., Fairbanks, H. R., Doostan, A., and Iaccarino, G., "Multi-Fidelity Uncertainty Quantification of Irradiated Particle-Laden Turbulence," Center for Turbulence Research TR SAND2017-12735B 659000, Stanford, CA, 2017, <https://www.osti.gov/biblio/1422682>.
- [30] Bachmayr, M., and Cohen, A., "Kolmogorov Widths and Low-Rank Approximations of Parametric Elliptic PDEs," *Mathematics of Computation*, Vol. 86, No. 304, 2017, pp. 701–724.
doi:10.1090/mcom/2017-86-304
- [31] Bui-Thanh, T., Willcox, K., and Ghattas, O., "Parametric Reduced-Order Models for Probabilistic Analysis of Unsteady Aerodynamic Applications," *AIAA Journal*, Vol. 46, No. 10, 2008, pp. 2520–2529.
doi:10.2514/1.35850
- [32] Narayan, A., Gittelsohn, C., and Xiu, D., "A Stochastic Collocation Algorithm with Multifidelity Models," *SIAM Journal on Scientific Computing*, Vol. 36, No. 2, Jan. 2014, pp. A495–A521.
doi:10.1137/130929461
- [33] Zhu, X., Narayan, A., and Xiu, D., "Computational Aspects of Stochastic Collocation with Multifidelity Models," *SIAM/ASA Journal on Uncertainty Quantification*, Vol. 2, No. 1, Jan. 2014, pp. 444–463.
doi:10.1137/130949154
- [34] Hampton, J., Fairbanks, H., Narayan, A., and Doostan, A., "Practical Error Bounds for a Non-intrusive Bi-fidelity Approach to Parametric/Stochastic Model Reduction," *Journal of Computational Physics*, Vol. 368, Sept. 2018, pp. 315–332, <https://www.sciencedirect.com/science/article/pii/S0021999118302298>
- [35] Spalart, P., and Allmaras, S., "A One-Equation Turbulence Model for Aerodynamic Flows," *30th Aerospace Sciences Meeting and Exhibit*, AIAA Paper 1992-439, 1992.
doi:10.2514/6.1992-439
- [36] Abbott, I. H., Von Doenhoff, A. E., and Stivers, L., "NACA Report 824: Summary of Airfoil Data," NACA TR-824, 1945, <https://ntrs.nasa.gov/search.jsp?R=19930090976>.
- [37] Carlson Biedron, R. T., "FUN3d Manual: 13.1," NASA TM-2017-219580, Feb. 2017, <https://ntrs.nasa.gov/search.jsp?R=20170002585>.
- [38] Whiting, C. H., and Jansen, K. E., "A Stabilized Finite Element Method for the Incompressible Navier-Stokes Equations Using a Hierarchical Basis," *International Journal for Numerical Methods in Fluids*, Vol. 35, No. 1, Jan. 2001, pp. 93–116.
doi:10.1002/(ISSN)1097-0363
- [39] Diskin, B., Thomas, J., Rumsey, C. L., and Schwoeppe, A., "Grid Convergence for Turbulent Flows," *53rd AIAA Aerospace Sciences Meeting, AIAA SciTech Forum*, AIAA Paper 2015-1746, 2015.
doi:10.2514/6.2015-1746
- [40] Streett Spalart, P. R., "Young-Person's Guide to Detached-Eddy Simulation Grids," NASA CR-2001-211032, July 2001, <https://ntrs.nasa.gov/search.jsp?R=20010080473>.
- [41] Freeman, J., and Roy, C., "Verification and Validation of Reynolds-Averaged Navier-Stokes Turbulence Models for External Flow," *Aerospace Science and Technology*, Vol. 32, No. 1, 2014, pp. 84–93, <https://www.sciencedirect.com/science/article/pii/S1270963813002022>.
- [42] Stein, K., Tezduyar, T., and Benney, R., "Mesh Moving Techniques for Fluid-Structure Interactions with Large Displacements," *Journal of Applied Mechanics*, Vol. 70, No. 1, Jan. 2003, pp. 58–63.
doi:10.1115/1.1530635
- [43] Ayachit, U., et al., "Performance Analysis, Design Considerations, and Applications of Extreme-Scale In Situ Infrastructures," *SC16: International Conference for High Performance Computing, Networking, Storage and Analysis*, IEEE Press, Piscataway, NJ, Nov. 2016, pp. 921–932.
doi:10.1109/SC.2016.78
- [44] Fabian, N., Moreland, K., Thompson, D., Bauer, A. C., Marion, P., Gevecik, B., Rasquin, M., and Jansen, K. E., "The Paraview Coprocessing Library: A Scalable, General Purpose In Situ Visualization Library," *2011 IEEE Symposium on Large Data Analysis and Visualization*, IEEE Publ., Piscataway, NJ, Oct. 2011, pp. 89–96.
doi:10.1109/LDAV.2011.6092322
- [45] Gu, M., and Eisenstat, S., "Efficient Algorithms for Computing a Strong Rank-Revealing QR Factorization," *SIAM Journal on Scientific Computing*, Vol. 17, No. 4, July 1996, pp. 848–869.
doi:10.1137/0917055
- [46] Cheng, H., Gimbutas, Z., Martinsson, P., and Rokhlin, V., "On the Compression of Low Rank Matrices," *SIAM Journal on Scientific Computing*, Vol. 26, No. 4, Jan. 2005, pp. 1389–1404.
doi:10.1137/030602678
- [47] Doostan, A., Geraci, G., and Iaccarino, G., "A Bi-Fidelity Approach for Uncertainty Quantification of Heat Transfer in a Rectangular Ribbed Channel," *Turbo Expo 2016: Turbomachinery Technical Conference and Exposition Volume 2C: Turbomachinery*, Paper GT2016-58092, Seoul, South Korea, June 2016, p. V02CT45A031.
doi:10.1115/GT2016-58092
- [48] Hastie, T., Tibshirani, R., and Friedman, J., *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., Springer–Verlag, New York, 2016.
- [49] Trefethen, L. N., and Bau, D., III, *Numerical Linear Algebra*, Soc. for Industrial and Applied Mathematics, Philadelphia, PA, June 1997, pp. 25–37.