

Multifidelity Monte Carlo estimation for large-scale uncertainty propagation*

Benjamin Peherstorfer[†]

University of Wisconsin-Madison, Madison, WI 53703

Philip S Beran[‡]

U.S. Air Force Research Laboratory, Wright-Patterson Air Force Base, Dayton, OH 45433

Karen Willcox[§]

Massachusetts Institute of Technology, Cambridge, MA 02139

One important task of uncertainty quantification is propagating input uncertainties through a system of interest to quantify the uncertainties' effects on the system outputs; however, numerical methods for uncertainty propagation are often based on Monte Carlo estimation, which can require large numbers of numerical simulations of the numerical model describing the system response to obtain estimates with acceptable accuracies. Thus, if the model is computationally expensive to evaluate, then Monte-Carlo-based uncertainty propagation methods can quickly become computationally intractable. We demonstrate that multifidelity methods can significantly speedup uncertainty propagation by leveraging low-cost low-fidelity models and establish accuracy guarantees by using occasional recourse to the expensive high-fidelity model. We focus on the multifidelity Monte Carlo method, which is a multifidelity approach that optimally distributes work among the models such that the mean-squared error of the multifidelity estimator is minimized for a given computational budget. The multifidelity Monte Carlo method is applicable to general types of low-fidelity models, including projection-based reduced models, data-fit surrogates, response surfaces, and simplified-physics models. We apply the multifidelity Monte Carlo method to a coupled aero-structural analysis of a wing and a flutter problem with a high-aspect-ratio wing. The low-fidelity models are data-fit surrogate models derived with standard procedures that are built in common software environments such as MATLAB and numpy/scipy. Our results demonstrate speedups of orders of magnitude compared to using the high-fidelity model alone.

I. Introduction

To meet the ever increasing demands on robustness and reliability of modern aircraft, it becomes necessary to take into account the effects of uncertainties on the system performance during design and operation. Uncertainty propagation is one task of uncertainty quantification that quantifies how uncertainties in the inputs of a system affect the uncertainties in the outputs of that system. If the system inputs are modeled as random variables, then the outputs of the system become random variables as well. The goal of uncertainty propagation is to characterize the output random variable, for example by estimating its statistics such as the expected value and the variance. Monte Carlo methods are widely used for uncertainty propagation. With a given numerical model that describes the input-output relationship of the system of interest, Monte Carlo methods evaluate the model at many input samples of the input random variable and estimate statistics from the corresponding outputs. Often a large number of model evaluations is required to estimate the statistics of the output random variable with an acceptable accuracy, see Figure 1a. If each evaluation of the model is computationally expensive, Monte Carlo estimation of statistics of the output random variable can become computationally intractable. The multifidelity Monte Carlo (MFMC) method leverages low-cost low-fidelity models of the high-fidelity model to speedup the estimation and occasionally uses recourse to the expensive high-fidelity model to establish unbiased estimators. In this paper, we demonstrate on a coupled aero-structural analysis of a wing

*Paper approved for public release (88ABW-2017-5909).

[†]Assistant Professor, Department of Mechanical Engineering and Wisconsin Institute for Discovery

[‡]Principal Research Aerospace Engineer, AFRL, AIAA Associate Fellow

[§]Professor, Department of Aeronautics and Astronautics, AIAA Associate Fellow




Fig. 1 Multifidelity methods combine high- and low-fidelity models. The low-fidelity models are leveraged for speedup and occasional recourse to the high-fidelity model is used to establish accuracy and/or convergence guarantees. This figure is a modified figure from the survey [16] on multifidelity methods.

and a flutter problem that MFMC can achieve speedups of orders of magnitude compared to Monte Carlo estimators that use the high-fidelity model alone.

Low-fidelity models provide approximations of the high-fidelity model outputs. The reduced accuracy compared to the high-fidelity model is in favor of significantly reduced evaluation costs. Low-fidelity models can be categorized into three different types. First, there are data-fit surrogate models [1, 2], e.g., kriging models [3–7], which are fitted to input-output data of the high-fidelity model. Second, there are projection-based reduced models [8–13] that solve the governing equations of the high-fidelity model in a problem-dependent reduced space, instead of the often high-dimensional general solution space. Third, simplified models [14, 15] are derived by simplifying the high-fidelity model, e.g., ignoring nonlinear effects and using a coarser grid to discretize the governing equations of the low-fidelity model. The survey [16] discusses these three low-fidelity model types in more detail.

A typical approach to speedup uncertainty propagation is to construct a low-fidelity model with one-time high costs and then use it to replace the high-fidelity model in the uncertainty propagation task. This means that the low-fidelity model is evaluated instead of the high-fidelity model, see Figure 1b. However, since the high-fidelity model is replaced by the low-fidelity model, the corresponding estimator is biased because the statistics of the low-fidelity model output random variable are estimated instead of the statistics of the high-fidelity model output random variable.

Multifidelity methods combine, instead of replace, the high-fidelity model with one or multiple low-fidelity models [16], see Figure 1c. The low-fidelity models are leveraged for speedup and occasional recourse to the high-fidelity model establishes accuracy guarantees. In case of uncertainty propagation, accuracy guarantees typically means unbiasedness with respect to the statistics of the high-fidelity model output random variable. Multifidelity methods rely on multiple models, use them in concert to speedup computations, and provide the same accuracy guarantees as methods that use the high-fidelity model alone. Many multifidelity methods can leverage heterogeneous types of low-fidelity models and so exploit the full breath of available low-fidelity models. For uncertainty propagation, there is a line of work on multifidelity methods [17–19] for stochastic collocation [20–22], which uses deterministic quadrature rules that can exploit, e.g., smoothness properties in the high- and low-fidelity model, instead of Monte Carlo methods, which are generally applicable but cannot exploit these problem structures. Multifidelity uncertainty propagation methods based on Monte Carlo often rely on control variates [23, 24] to obtain a multifidelity estimator with a lower variance than a single-fidelity Monte Carlo estimator of the same costs. The work [25, 26] uses reduced-basis models as control variates to reduce the variance. The StackMC method [27] relies on data-fit surrogate models to construct control variates for variance reduction. The multilevel Monte Carlo method [28, 29] exploits a hierarchy of coarse-grid approximations of the governing equations of the high-fidelity model to construct control variates. An extension of the multilevel Monte Carlo method to reduced-basis models is introduced in the work [30].

We focus here on the MFMC method [31, 32] that uses control variates to reduce the variance in Monte Carlo




Fig. 2 Uncertainty propagation is the task of propagating input uncertainties through a model to characterize the uncertainties' effects on the model outputs.

estimation and relies on low-fidelity models for constructing these control variates. MFMC optimally decides how often each of the models has to be evaluated to minimize the mean-squared error (MSE) of the MFMC estimator for a given computational budget. We demonstrate MFMC on a coupled aero-structural analysis and a flutter problem. The low-fidelity models are derived with standard procedures available in MATLAB. MFMC achieves speedups of orders of magnitude in our results compared to using Monte Carlo estimators with the high-fidelity model alone.

II. Uncertainty propagation with Monte Carlo estimation

This section sets up the problem and discusses Monte Carlo estimation for uncertainty propagation.

A. Problem setup

Consider a high-fidelity model that describes the response of a system of interest for given inputs. The inputs are the components of a d -dimensional vector $\mathbf{x} = [x_1, \dots, x_d]^T$ in an input domain $\mathcal{X} \subseteq \mathbb{R}^d$ with $d \in \mathbb{N}$. Evaluating the high-fidelity model at an input \mathbf{x} gives an approximation $y \in \mathcal{Y}$ of the system response in the output domain $\mathcal{Y} \subseteq \mathbb{R}$. The output y is our quantity of interest. We denote the high-fidelity model as a function

$$f : \mathcal{X} \rightarrow \mathcal{Y}$$

that maps inputs $\mathbf{x} \in \mathcal{X}$ onto outputs $y = f(\mathbf{x})$. In the following, we consider the case where f is computationally expensive to evaluate. Each evaluation of f incurs costs $w \in \mathbb{R}$ with $w > 0$.

Consider now a random variable $X : \Omega \rightarrow \mathcal{X}$ with sample space Ω . The random variable X describes the inputs and accounts for the uncertainties in the inputs. Since X is a random variable, the output $f(X)$ becomes a random variable as well. We are interested in estimating statistics of $f(X)$ for a given model f and a given input random variable X . Figure 2 depicts this uncertainty propagation task. In the following, we restrict the discussion to estimating the expected value

$$s = \mathbb{E}[f(X)] \tag{1}$$

of the output random variable $f(X)$. The expected value is

$$\mathbb{E}[f(X)] = \int_{\mathcal{X}} f(\mathbf{x})p(\mathbf{x})d\mathbf{x},$$

with the probability density function p corresponding to the input random variable X .

B. Monte Carlo estimation

A Monte Carlo estimator of the expected value s of $f(X)$ is

$$\bar{y}_m = \frac{1}{m} \sum_{i=1}^m f(\mathbf{x}_i), \tag{2}$$

where $\mathbf{x}_1, \dots, \mathbf{x}_m$ are m realizations of the random variable X . The Monte Carlo estimator \bar{y}_m is an unbiased estimator of s because

$$\mathbb{E}[\bar{y}_m] = s.$$

The MSE of \bar{y}_m is defined as

$$e(\bar{y}_m) = \mathbb{E} \left[(\bar{y}_m - s)^2 \right],$$

which simplifies to

$$e(\bar{y}_m) = \frac{\text{Var}[f(X)]}{m},$$

because \bar{y}_m is unbiased. Thus, the MSE of the Monte Carlo estimator \bar{y}_m depends on the variance $\text{Var}[f(X)]$ and on the number of samples m . The costs of the Monte Carlo estimator \bar{y}_m are

$$c(\bar{y}_m) = wm,$$

because one evaluation of the high-fidelity model f incurs costs w and the Monte Carlo estimator evaluates f at m realizations of X . Note that only the evaluation costs of f are taken into account. Other costs, e.g., computing the sum in (2) once the high-fidelity model f is evaluated at the realizations $\mathbf{x}_1, \dots, \mathbf{x}_m$, are typically negligible.

Depending on the variance $\text{Var}[f(X)]$ of the output random variable, a large number of realizations, and thus a large number of high-fidelity model evaluations, can be necessary to obtain a Monte Carlo estimate of s with an acceptable MSE. If the high-fidelity model is expensive to evaluate, i.e., if w is large, then Monte Carlo estimation without further ado can become computationally intractable quickly.

III. Multifidelity Monte Carlo estimation for uncertainty propagation

The multifidelity Monte Carlo [31, 32] (MFMC) method aims to leverage low-cost low-fidelity models to obtain a multifidelity estimator that achieves the same MSE as the standard Monte Carlo estimator (2) with significantly reduced costs. In a typical situation, large numbers of evaluations of the low-cost low-fidelity models are taken and only a few evaluations of the high-fidelity model. The large numbers of low-fidelity model evaluations help to reduce the variance of the multifidelity estimator whereas the few and occasional evaluations of the high-fidelity model establish unbiasedness of the multifidelity estimator. We first introduce the notation for low-fidelity models in Section III-A and then discuss MFMC in Section III-B. Implementation details of MFMC are given in Section III-C.

A. Low-fidelity models

In the following, we denote the high-fidelity model as $f^{(1)}$, i.e., with a superscript (1), and the low-fidelity models as $f^{(2)}, \dots, f^{(k)}$ with $k \in \mathbb{R}$. Thus, we have one high-fidelity model $f^{(1)}$ and $k - 1$ low-fidelity models $f^{(2)}, \dots, f^{(k)}$. All $i = 2, \dots, k$ low-fidelity models $f^{(i)} : \mathcal{X} \rightarrow \mathcal{Y}$ map from the same input domain \mathcal{X} to the same output domain \mathcal{Y} as the high-fidelity model and approximate the high-fidelity model output, see Figure 3. Evaluating a model $f^{(i)}$ incurs costs w_i with $w_i \in \mathbb{R}$ and $0 < w_i$, for $i = 1, \dots, k$. No assumptions are made on the low-fidelity models in the following. In particular, we do not make assumptions on the type of the low-fidelity model, the costs, and their approximation quality with respect to the high-fidelity model.

B. Multifidelity Monte Carlo estimation

Consider evaluating the high-fidelity model $f^{(1)}$ m_1 times, low-fidelity model $f^{(2)}$ m_2 times, and so on, and evaluating low-fidelity model $f^{(k)}$ m_k times. In the following, we have $0 < m_1 \leq m_2 \leq \dots \leq m_k$, which means that the high-fidelity model m_1 is evaluated at least once and fewer or equal many times as the low-fidelity model $f^{(2)}$, and so on. The low-fidelity model $f^{(k)}$ is evaluated the most times. Consider now m_k realizations

$$\mathbf{x}_1, \dots, \mathbf{x}_{m_k} \tag{3}$$

of the random variable X . Evaluate model $f^{(i)}$ at the first m_i realizations $\mathbf{x}_1, \dots, \mathbf{x}_{m_i}$ of the realizations (3) to obtain the outputs

$$f^{(i)}(\mathbf{x}_1), \dots, f^{(i)}(\mathbf{x}_{m_i}), \tag{4}$$




Fig. 3 The MFMC method leverages the low-fidelity models $f^{(2)}, \dots, f^{(k)}$ to speedup uncertainty propagation and uses occasional recourse to the high-fidelity model $f^{(1)}$ to guarantee unbiased estimators.

for $i = 1, \dots, k$. Thus, the high-fidelity model $f^{(1)}$ is evaluated at the first m_1 realizations $\mathbf{x}_1, \dots, \mathbf{x}_{m_1}$, low-fidelity model $f^{(2)}$ is evaluated at the realizations $\mathbf{x}_1, \dots, \mathbf{x}_{m_2}$, and so on. With the outputs (4), define the Monte Carlo estimators

$$\bar{y}_{m_1} = \frac{1}{m_1} \sum_{j=1}^{m_1} f^{(1)}(\mathbf{x}_j), \quad \bar{y}_{m_i}^{(i)} = \frac{1}{m_i} \sum_{j=1}^{m_i} f^{(i)}(\mathbf{x}_j), \quad \bar{y}_{m_{i-1}}^{(i)} = \frac{1}{m_{i-1}} \sum_{j=1}^{m_{i-1}} f^{(i)}(\mathbf{x}_j), \quad (5)$$

for $i = 2, \dots, k$. The Monte Carlo estimator \bar{y}_{m_1} is derived from high-fidelity model outputs, and the Monte Carlo estimators $\bar{y}_{m_i}^{(i)}$ and $\bar{y}_{m_{i-1}}^{(i)}$ are derived from low-fidelity model outputs. The estimator $\bar{y}_{m_i}^{(i)}$ uses all evaluations $f^{(i)}(\mathbf{x}_1), \dots, f^{(i)}(\mathbf{x}_{m_i})$ of the low-fidelity model $f^{(i)}$, whereas the estimator $\bar{y}_{m_{i-1}}^{(i)}$ uses the first m_{i-1} evaluations of the low-fidelity model $f^{(i)}$ only.

The MFMC method combines the Monte Carlo estimators (5) into the MFMC estimator

$$\hat{s} = \bar{y}_{m_1} + \sum_{i=2}^k \alpha_i \left(\bar{y}_{m_i}^{(i)} - \bar{y}_{m_{i-1}}^{(i)} \right). \quad (6)$$

The MFMC estimator depends on the coefficients $\alpha_2, \dots, \alpha_k \in \mathbb{R}$ and on the number of model evaluations m_1, \dots, m_k . Before we can specify how to select $m_1, \dots, m_k, \alpha_2, \dots, \alpha_k$, we first have to define the costs of the MFMC estimator. The costs of \hat{s} are

$$c(\hat{s}) = \sum_{i=1}^k m_i w_i, \quad (7)$$

because model $f^{(i)}$ is evaluated m_i times and each evaluation incurs costs w_i for $i = 1, \dots, k$. The MFMC method selects $m_1, \dots, m_k, \alpha_2, \dots, \alpha_k$ such that the costs of the MFMC estimator are equal to a given computational budget $c(\hat{s}) = p$ and the MSE of the MFMC estimator is minimal. Thus, the MFMC estimator is optimal in the sense that $m_1, \dots, m_k, \alpha_2, \dots, \alpha_k$ are selected that minimize the MSE while the costs of the MFMC estimator equal the given computational budget p .

Key to optimally select the $m_1, \dots, m_k, \alpha_2, \dots, \alpha_k$ are the correlation coefficients between the low-fidelity models and the high-fidelity model. To define the correlation coefficient, let Z_1 and Z_2 be two random variables. The correlation coefficient of Z_1 and Z_2 is

$$\rho = \frac{\text{Cov}[Z_1, Z_2]}{\sqrt{\text{Var}[Z_1] \text{Var}[Z_2]}},$$

where

$$\text{Cov}[Z_1, Z_2] = \mathbb{E}[(Z_1 - \mathbb{E}[Z_1])(Z_2 - \mathbb{E}[Z_2])]$$

is the covariance of Z_1 and Z_2 . The correlation coefficient is normalized such that $-1 \leq \rho \leq 1$. If $|\rho| \approx 1$, then Z_1 and Z_2 are highly correlated, which means that Z_1 and Z_2 behave similarly. If $|\rho| \approx 0$, then Z_1 and Z_2 are weakly correlated, and the behavior of Z_1 gives only little information about the behavior of Z_2 and vice versa. The MFMC method uses the correlation coefficients between the high-fidelity model output random variable $f^{(1)}(X)$ and the low-fidelity model output random variables $f^{(i)}(X)$

$$\rho_i = \frac{\text{Cov}[f^{(1)}(X), f^{(i)}(X)]}{\sqrt{\text{Var}[f^{(1)}(X)] \text{Var}[f^{(i)}(X)]}}, \quad i = 2, \dots, k, \quad (8)$$

and the costs w_1, \dots, w_k of the models to optimally select the number of model evaluations m_1, \dots, m_k and the coefficients $\alpha_2, \dots, \alpha_k$. If a low-fidelity model $f^{(i)}$ is cheap to evaluate, i.e., the costs of the low-fidelity model are significantly lower than the costs of the high-fidelity model $w_i \ll w_1$, and the low-fidelity model output random variable $f^{(i)}(X)$ is highly correlated to the high-fidelity model output random variable $f^{(1)}(X)$, then the low-fidelity model provides valuable information about the high-fidelity model at low costs. Thus, in this situation, MFMC tends to evaluate this low-fidelity model many times. In contrast, if a low-fidelity model is computationally expensive and weakly correlated, then MFMC tends to evaluate the low-fidelity model only a few times or does not evaluate it at all. Expressions for the optimal $m_1, \dots, m_k, \alpha_2, \dots, \alpha_k$ with respect to the costs w_1, \dots, w_k and the correlation coefficients ρ_2, \dots, ρ_k are available [32].

C. Implementation of MFMC estimation

A detailed algorithm that implements MFMC estimation is given in the work [32]. A MATLAB implementation of MFMC is available on GitHub*. The MFMC implementation requires estimates of the costs w_1, \dots, w_k and of the correlation coefficients ρ_2, \dots, ρ_k of the high-fidelity model output random variable and the low-fidelity model output random variables, see Section III-B and equation (8). All these quantities can be estimated from pilot runs with the models $f^{(1)}, \dots, f^{(k)}$, as described in detail in [32, Section 3.6]. The estimates of the costs and the correlation coefficients are required to optimally select the number of model evaluations m_1, \dots, m_k and the coefficients $\alpha_2, \dots, \alpha_k$. Errors in the estimates of the costs and the correlation coefficients affect the choice of $m_1, \dots, m_k, \alpha_2, \dots, \alpha_k$ only and cannot introduce a bias into the MFMC estimator of s . Thus, crude estimates of the costs and correlation coefficients are typically sufficient, cf. the discussion in [32, Section 3.4].

IV. Numerical experiments

We demonstrate MFMC on two examples. In Section IV-A, we use MFMC to propagate uncertain environment conditions through a coupled aero-structural wing analysis to quantify the uncertainty in the fuel burn. In Section IV-B and Section IV-C, we consider a flutter problem with a high-aspect-ratio wing. All runtime measurements are performed with MATLAB and python on compute nodes with Intel Xeon E5-1660v4 and 64GB RAM.

A. Uncertainty propagation in a coupled aero-structural analysis

This sections uses MFMC to propagate uncertain flight conditions and uncertain environmental conditions through a coupled aero-structural analysis of a wing.

1. Problem setup

The high-fidelity model corresponds to a coupled aero-structural analysis code called OpenAeroStruct† [33] that utilizes a vortex-lattice method and a 6 degree of freedom 3-dimensional spatial beam model. The code simulates aerodynamic and structure analysis using lifting surfaces. The code is built with the framework OpenMDAO [34], which is a platform for systems analyses and multidisciplinary optimization. An illustration of a wing analyzed with the high-fidelity model is given in Figure 4.

We consider wing meshes of 5 evenly spaced spanwise and 2 chordwise points, which is the default configuration in the code. The input is a three-dimensional vector $\mathbf{x} = [x_1, x_2, x_3]^T$, with x_1 describing the angle of attack (deg), x_2 air density (kg/m^3), and x_3 Mach number. The rest of the parameters of OpenAeroStruct are set to their default values. The output is fuel burn. The input domain is $\mathcal{X} = \mathbb{R}^3$ and the output domain is $\mathcal{Y} = \mathbb{R}$. The high-fidelity model $f^{(1)} : \mathcal{X} \rightarrow \mathcal{Y}$ uses OpenAeroStruct to map a realization \mathbf{x} of the input random variable X onto fuel burn.

To model uncertainties in the inputs, we consider the input random variable X . The distribution of X is a mixture distribution of two normal distributions with mean $\boldsymbol{\mu}_1 = [5, 0.38, 0.80]^T$ and covariance

$$\boldsymbol{\Sigma}_1 = \begin{bmatrix} 0.01 & 0 & 0 \\ 0 & 0.001 & 0 \\ 0 & 0 & 0.001 \end{bmatrix}$$

and mean $\boldsymbol{\mu}_2 = [5, 0.38, 0.84]^T$ and covariance

$$\boldsymbol{\Sigma}_2 = \begin{bmatrix} 0.01 & 0 & 0 \\ 0 & 0.001 & 0 \\ 0 & 0 & 0.006 \end{bmatrix}.$$

With the mixture distribution, we model small variations about the base values $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$, in contrast to variations in intervals, as with uniform distributions. The base values $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$ are given as examples in the OpenAeroStruct code. Note that MFMC is applicable to any input distribution that satisfies certain mild assumptions, e.g., that $\mathbb{E}[X]$ and $\text{Var}[X]$ exist in \mathbb{R} , see [32] for details. In particular, the components of the input random variable can be correlated. A realization of X is an input \mathbf{x} that can be used to evaluate the high-fidelity model to obtain the corresponding fuel burn. Our goal is to estimate the expected fuel burn $\mathbb{E}[f^{(1)}(X)]$.

*<https://github.com/pehersto/mfmc/>

†<https://github.com/johnjasa/OpenAeroStruct/>




Fig. 4 Illustration of a wing that can be analyzed with the OpenAeroStruct code. Figure from OpenAeroStruct documentation[†].

Table 1 Coupled aero-structural analysis: The high-fidelity model $f^{(1)}$ is six orders of magnitude more expensive to evaluate than the low-fidelity model $f^{(2)}$ and $f^{(3)}$ in this example. The costs of constructing (“offline costs”) the low-fidelity model $f^{(2)}$ are almost three times higher than the costs of constructing the low-fidelity model $f^{(3)}$.

model	evaluation costs [s]	offline costs [s]	correlation coefficient
high-fidelity model $f^{(1)}$	1.61×10^{-1}	-	-
low-fidelity model $f^{(2)}$, $n = 7$	1.23×10^{-7}	55.382	9.9552×10^{-1}
low-fidelity model $f^{(3)}$, $n = 5$	1.21×10^{-7}	20.183	9.9192×10^{-1}

2. Low-fidelity models

Spline interpolants of the high-fidelity model $f^{(1)}$ serve as low-fidelity models. Let $n \in \mathbb{N}$ and consider an equidistant grid with n grid points in each direction in the domain

$$[4.5, 5.5] \times [0.2, 0.6] \times [0.3, 1.25] \subset \mathbb{R}^3. \quad (9)$$

Note that the OpenAeroStruct code uses a linear method for the aerodynamic analysis and therefore fails to capture the behavior in the transonic regime. The distribution of the input random variable X gives realizations that correspond to the transonic regime with a very low probability. To avoid numerical problems with outliers that fall into the transonic regime, the domain for constructing the low-fidelity models (9) includes Mach numbers $x_3 > 1$.

To construct a spline interpolant, the high-fidelity model $f^{(1)}$ is evaluated at all grid points and the interpolant is constructed with the `griddedInterpolant` that is built in MATLAB. The costs of constructing a spline interpolant are $n^3 \times w_1$, because n^3 evaluations of the high-fidelity model are required and each evaluation of the high-fidelity model incurs costs w_1 .

We consider two low-fidelity models: The low-fidelity $f^{(2)}$ is constructed with n set to $n = 7$, i.e., from $n^3 = 7^3 = 343$ outputs of the high-fidelity model. The low-fidelity model $f^{(3)}$ is constructed with $n = 5$. The properties of the low- and the high-fidelity models are summarized in Table 1. The costs and the correlation coefficients are estimated from 1000 samples following the procedure described in [32, Section 3.6].

3. Results

We estimate $\mathbb{E}[f^{(1)}(X)]$ with MFMC using the models with properties reported in Table 1. We first compute a reference estimate of $\mathbb{E}[f^{(1)}(X)]$ using MFMC and a computational budget that is equal to the costs of 10^4 evaluations of the high-fidelity model. This computation is repeated 50 times and the average of the corresponding estimates is our

Table 2 Coupled aero-structural analysis: Table reports the number of times the low- and the high-fidelity models are evaluated in the Monte Carlo and the MFMC estimators. For example, for online costs of about 800 seconds, the Monte Carlo estimator evaluates the high-fidelity model 5000 times, whereas the MFMC estimators split these 800 seconds between evaluating the low- and high-fidelity models. In case of the MFMC estimator that uses all three models $f^{(1)}, f^{(2)}, f^{(3)}$, the high-fidelity model is evaluated about 4950 times and the remaining 50 high-fidelity model evaluations are traded for millions of evaluations of the low-fidelity models.

online costs [s]	Monte Carlo	MFMC with $f^{(1)}, f^{(2)}$		MFMC with $f^{(1)}, f^{(3)}$		MFMC with $f^{(1)}, f^{(2)}, f^{(3)}$		
	#evals $f^{(1)}$	#evals $f^{(1)}$	#evals $f^{(2)}$	#evals $f^{(1)}$	#evals $f^{(3)}$	#evals $f^{(1)}$	#evals $f^{(2)}$	#evals $f^{(3)}$
7.99×10^0	50	4.90×10^1	4.48×10^5	4.90×10^1	5.97×10^5	4.90×10^1	5.07×10^4	5.99×10^5
1.61×10^1	100	9.90×10^1	8.95×10^5	9.90×10^1	1.19×10^6	9.90×10^1	1.01×10^5	1.20×10^6
8.07×10^1	500	4.96×10^2	4.48×10^6	4.95×10^2	5.97×10^6	4.95×10^2	5.07×10^5	5.99×10^6
1.61×10^2	1000	9.93×10^2	8.95×10^6	9.90×10^2	1.19×10^7	9.90×10^2	1.01×10^6	1.20×10^7
8.07×10^2	5000	4.97×10^3	4.48×10^7	4.95×10^3	5.97×10^7	4.95×10^3	5.07×10^6	5.99×10^7

reference estimate \hat{s}^{Ref} . Note that the MFMC estimator is guaranteed to be unbiased (and consistent) and therefore using MFMC with a large computational budget to compute the reference estimate is reasonable. We then estimate $\mathbb{E}[f^{(1)}(X)]$ with three different estimators: The single-fidelity Monte Carlo estimator \hat{s}^{HF} that uses the high-fidelity model $f^{(1)}$ alone. The single-fidelity Monte Carlo estimator \hat{s}^{LF} that uses the low-fidelity model $f^{(3)}$ alone. The MFMC estimator \hat{s}^{MFMC} that combines the high-fidelity model $f^{(1)}$ and the low-fidelity model $f^{(3)}$. We repeat the estimation 100 times to obtain the estimates $\hat{s}_1^{\text{HF}}, \dots, \hat{s}_{100}^{\text{HF}}$ from the high-fidelity model alone, the estimates $\hat{s}_1^{\text{LF}}, \dots, \hat{s}_{100}^{\text{LF}}$ from the low-fidelity model alone, and the MFMC estimates $\hat{s}_1^{\text{MFMC}}, \dots, \hat{s}_{100}^{\text{MFMC}}$. The estimated relative MSE of the corresponding estimators is

$$\hat{e}(\hat{s}) = \frac{1}{100} \sum_{i=1}^{100} \frac{(\hat{s}_i - \hat{s}^{\text{Ref}})^2}{(\hat{s}^{\text{Ref}})^2},$$

where \hat{s} stands for either $\hat{s}^{\text{HF}}, \hat{s}^{\text{LF}}$, or \hat{s}^{MFMC} .

Figure 5a reports the estimated relative MSE of the two single-fidelity and the multifidelity estimator for different online costs. The online costs equal the computational budget p that is available for evaluating the high-fidelity and the low-fidelity models, see equation (7). It includes the evaluation costs only and ignores the costs of constructing the low-fidelity models. Furthermore, the online costs are with respect to one estimation, instead of the sum of all 100 estimations over which the relative MSE is estimated. The estimated relative MSE corresponding to the single-fidelity estimator that uses the low-fidelity model alone levels off, which demonstrates that using the low-fidelity model alone introduces a bias. For the same online costs, the MFMC estimator achieves an estimated relative MSE that is more than one order of magnitude lower than the single-fidelity Monte Carlo estimator that uses the high-fidelity model alone. The estimated relative MSE of the MFMC estimator is reduced as the online costs are increased, which indicates that the MFMC estimator is unbiased, even though the low-fidelity model is used to obtain speedups compared to using the high-fidelity model alone. Figure 5b shows the estimated relative MSE corresponding to the three MFMC estimators that use different combinations of the three models $f^{(1)}, f^{(2)}, f^{(3)}$. The MFMC estimator that uses $f^{(1)}$ and $f^{(2)}$ achieves a comparable estimated MSE error as the MFMC estimator that uses $f^{(1)}$ and $f^{(3)}$, which is in agreement with Table 1 that shows that both low-fidelity models have similar properties in terms of costs and correlation coefficients. Combining all three models $f^{(1)}, f^{(2)}, f^{(3)}$ seems to give a small improvement compared to using two models only, which again can be explained with the similar properties of the two low-fidelity models.

Table 2 reports the number of evaluations of the low- and high-fidelity models for different online costs and estimators. The MFMC estimators evaluate the low-fidelity models orders of magnitude more often than the high-fidelity model. For example, for online costs of about 800 seconds, the high-fidelity model can be evaluated 5000 times, which means that the single-fidelity Monte Carlo estimator that uses the high-fidelity model alone is derived from 5000 realizations of $f^{(1)}(X)$. In contrast, for online costs of about 800 seconds, the MFMC estimator that uses the low-fidelity models $f^{(2)}, f^{(3)}$ and the high-fidelity model $f^{(1)}$, evaluates $f^{(1)}$ about 4950 times, model $f^{(2)}$ more than five million times, and model $f^{(3)}$ almost 60 million times.

Figure 6 reports the estimated relative MSE of the estimator with respect to the total costs, which includes the online costs of evaluating the models and the offline costs of constructing the low-fidelity models. Since the offline costs are small compared to the online costs in this example, the MFMC estimator achieves similar speedups with respect to the total costs as with respect to the online costs.




Fig. 5 Coupled aero-structural analysis: The multifidelity approach achieves a speedup of about two orders of magnitude. The online costs are the costs of evaluating the high-fidelity and the low-fidelity models. For example, for online costs of about 800 seconds, the high-fidelity model can be evaluated 5000 times, cf. Table 1. Note that the error curve in plot (a) corresponding to the single-fidelity approach that uses the low-fidelity model $f^{(3)}$ alone levels off because using the low-fidelity model alone introduces a bias into the estimator. The multifidelity approach avoids that bias by combining, instead of replacing, the high-fidelity model with the low-fidelity model.

B. Uncertainty propagation of flutter of a high-aspect-ratio wing

This numerical experiment studies the flutter of a highly flexible, high-aspect-ratio wing, see Figure 7. The air density and the root angle of attack are uncertain inputs and the expected value of the flutter speed is the statistic of the output random variable that is to be estimated.

1. Problem setup

The flutter problem that we consider was studied experimentally and theoretically by Tang and Dowell [36]. The aeroelastic response of the wing is analyzed in the context of a structural design optimization method in [35] and a sensitivity analysis of the flutter speed of this wing is carried out in [37]. In [37], flutter speed and analytical sensitivities of flutter speed with respect to input parameters are computed to near machine precision with regards to temporal discretization, thus removing this source of error from the Monte Carlo and MFMC estimates reported below. For the wing model and the parameter space studied here, the flutter solutions appear unique and well behaved. The consistency of the flutter solutions with time-domain analysis has previously been established by the second author. We refer to [38] for a detailed discussion on aeroelasticity and uncertainty quantification, where the authors establish that non-uniqueness in the flutter solution can be modeled as bi-modal distributions.

We briefly summarize the wing model here and refer to the literature [35–37] for details. The elastic and geometrically nonlinear behavior of the wing is modeled by the Hodges-Dowell equations [35]. Nonlinear terms of third order or higher are neglected. The structural equations are discretized [35] using $l \in \mathbb{N}$ finite elements on an equidistant grid from wing root to tip. The aerodynamic model is the ONERA stall model of Tran and Petot [39]. The resulting aerodynamical model is composed of $6(l + 1)$ degrees of freedom, which leads to $10l + 6(l + 1)$ equations if coupled with the structural equations in first-order form. We consider two inputs to the high-fidelity model, which are the root angle of attack x_1 (rad/s) and the air density x_2 (kg/m³). The input domain is

$$\mathcal{X} = [0.0332, 0.0367] \times [1.1638, 1.2863].$$

The inputs vary about their base value 0.0349 (about 2 degrees) and 1.225, respectively, by up to 5%. The output domain is $\mathcal{Y} = \mathbb{R}$ and the output is the flutter speed. We set the number of finite elements to $l = 10$ and define the high-fidelity model $f^{(1)} : \mathcal{X} \rightarrow \mathcal{Y}$ that maps the root angle of attack and the air density onto the flutter speed. To account for uncertainties in the root angle of attack and the air density, we consider the input random variable X that follows a uniform distribution in the input domain \mathcal{X} . The goal is to estimate the expected flutter speed $\mathbb{E}[f^{(1)}(X)]$.




Fig. 6 Coupled aero-structural analysis: The plots report the estimated relative MSE with respect to the total costs that include the construction of the low-fidelity models and the model evaluation costs, in contrast to Figure 5 that reports the MSE with respect to the online costs that include the evaluation costs of the models only. The reported results indicate that the costs of constructing the low-fidelity models are small compared to the online costs in this example.




Fig. 7 Schematic of model of high-aspect-ratio wing. In the figure, U denotes air speed and x, y, z the coordinate directions. Figure reproduced from the work [35].

2. Low-fidelity models

Low-fidelity models are derived via spline interpolation from the high-fidelity model $f^{(1)}$. Let $n \in \mathbb{N}$ be the number of equidistant grid points in each direction of the input domain \mathcal{X} . A spline interpolant is constructed with the MATLAB function `griddedInterpolant`. We consider the low-fidelity model $f^{(2)}$ that is derived from $n = 2$ grid points in each direction. Since we are interested in small variations of the root angle of attack and the air density only, and the problem seems insensitive to small perturbations [37], a few interpolation points are sufficient to obtain a low-fidelity model with an acceptable accuracy for MFMC estimation. The properties of the high- and low-fidelity model are summarized in Table 3. The costs and correlation coefficients are estimated as in Section IV-A-2.

While the flutter analysis procedure provides sensitivities with respect to the problem parameters, these sensitivities are not exploited in the MFMC method discussed herein. There is potential for the use of this sensitivity data to improve the quality of the low-fidelity models (e.g., gradient enhance kriging) or by using the linearization itself as a low-fidelity model.

3. Results

We estimate $\mathbb{E}[f^{(1)}(X)]$ with MFMC and standard Monte Carlo estimators and compare the estimated relative MSEs. We follow the same procedure as described in Section IV-A-3. The reference estimate \hat{y}^{Ref} is obtained with MFMC

Table 3 Flutter problem: The table reports the costs and correlation coefficients of the low- and high-fidelity model.

model	evaluation costs [s]	offline costs [s]	correlation coefficient
high-fidelity model $f^{(1)}$	1.51×10^1	-	-
low-fidelity model $f^{(2)}$, $n = 2$	5.39×10^{-8}	60.517	9.9981×10^{-1}

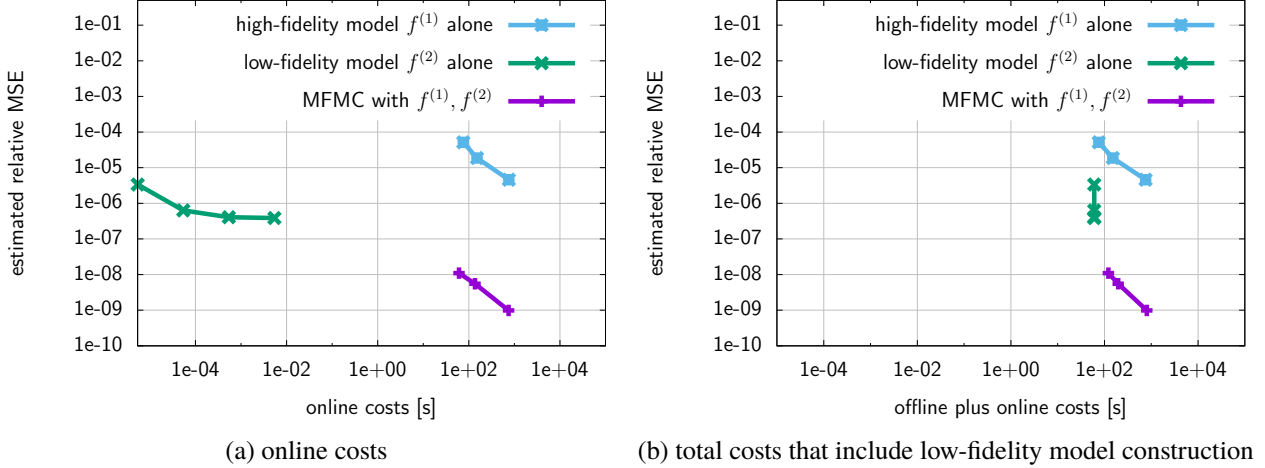


Fig. 8 Flutter problem: Using the low-fidelity model alone to speedup uncertainty propagation typically leads to biased estimators, as indicated by the leveling off of the error curve corresponding to the Monte Carlo estimator that uses the low-fidelity model alone. The MFMC approach leverages the low-fidelity model to obtain about three orders of magnitude speedup compared to using the high-fidelity model alone and guarantees an unbiased estimator.

using the high-fidelity model $f^{(1)}$ and the low-fidelity model $f^{(2)}$. The budget is equal to the costs of evaluating the high-fidelity model 1000 times. The MFMC estimation is repeated 100 times and the average estimate is our reference estimate \hat{s}^{Ref} . Figure 8 compares the estimated relative MSE of the single-fidelity Monte Carlo estimator that uses the high-fidelity model alone, the single-fidelity Monte Carlo estimator that uses the low-fidelity model alone, and the MFMC estimator that combines the low-fidelity with the high-fidelity model. The relative MSEs are estimated over 100 runs, as in Section IV-A-3. The results reported in Figure 8 show that the estimated relative MSE of the single-fidelity Monte Carlo estimator that uses the low-fidelity model alone levels off. The MFMC estimator leverages the low-fidelity model to obtain about three orders of magnitude speedup compared to the single-fidelity Monte Carlo estimator that uses the high-fidelity model alone and guarantees unbiasedness by occasional recourse to the high-fidelity model. The number of model evaluations of the estimators are summarized in Table 4. The MFMC estimator evaluates the low-fidelity model six orders of magnitude more often than the high-fidelity model.

Table 4 Flutter problem: The table reports the number of times the low- and the high-fidelity models are evaluated in the MFMC estimator and compares the number of high-fidelity model evaluations to the number of high-fidelity model evaluations in the standard Monte Carlo estimator that uses the high-fidelity model alone.

online costs [s]	Monte Carlo	MFMC with $f^{(1)}, f^{(2)}$	
	#evals $f^{(1)}$	#evals $f^{(1)}$	#evals $f^{(2)}$
7.56×10^1	5	4.00×10^0	4.30×10^6
1.51×10^2	10	9.00×10^0	8.61×10^6
7.56×10^2	50	4.90×10^1	4.30×10^7

Table 5 Flutter problem with three inputs: The table reports the costs and correlation coefficients of the low- and high-fidelity models. The output random variable corresponding to low-fidelity model $f^{(2)}$ is significantly higher correlated to the high-fidelity output random variable than the output random variable corresponding to model $f^{(3)}$.

model	evaluation costs [s]	offline costs [s]	correlation coefficient
high-fidelity model $f^{(1)}$	1.51×10^1	-	-
low-fidelity model $f^{(2)}, n = 7$	1.29×10^{-7}	4923	$9.9999997323 \times 10^{-1}$
low-fidelity model $f^{(3)}, n = 3$	1.18×10^{-7}	387.53	$9.9987563563 \times 10^{-1}$

C. Uncertainty propagation of flutter of a high-aspect-ratio wing with three inputs and strong variations

This section considers the same flutter problem as studied in Section IV-B but with additional inputs and stronger variations in the inputs.

1. Problem setup and low-fidelity models

The flutter problem of Section IV-B takes the root angle of attack x_1 and air density x_2 as inputs. We now consider the same flutter problem but with the mass of the tip x_3 (kg/m³) as an additional input. The input domain is

$$\mathcal{X} = [8.72 \times 10^{-3}, 4.3 \times 10^{-2}] \times [1.1638, 1.2863] \times [3.9615 \times 10^{-2}, 4.3785 \times 10^{-2}],$$

such that the variation of root angle of attack is from about 0.5 degrees to about 2.5 degrees. Note that this is a significantly stronger variation than in the flutter problem studied in Section IV-B. The air density and mass of the tip vary 5% around their base value 1.225 and 4.17×10^{-2} , respectively. Detailed description of these inputs and their effect on the flutter speed can be found in [37]. The high-fidelity model $f^{(1)} : \mathcal{X} \rightarrow \mathcal{Y}$ is derived with 10 finite elements, see Section IV-B, and maps the inputs onto the flutter speed. The input random variable X has independent components and is uniformly distributed in \mathcal{X} . Our goal is to estimate the expected flutter speed $\mathbb{E}[f^{(1)}(X)]$.

We construct two low-fidelity models with spline interpolation as discussed in Section IV-B-2. The low-fidelity model $f^{(2)}$ is derived from $n = 7$ grid points in each direction of the input domain \mathcal{X} , and the low-fidelity model $f^{(3)}$ is derived from $n = 3$ grid points in each direction. The properties of the high- and low-fidelity model are summarized in Table 5.

2. Results

Figure 9 reports the estimated relative MSEs of single-fidelity Monte Carlo and MFMC estimators. The reference estimate \hat{s}^{Ref} is the average over 100 estimates obtained with MFMC with models $f^{(1)}, f^{(2)}, f^{(3)}$ and a budget that is equal to the costs of 1000 high-fidelity model evaluations. The relative MSEs of the estimators in Figure 9 are obtained over 100 runs as in Section IV-A and Section IV-B. Consider first the MFMC estimator that uses models $f^{(1)}$ and $f^{(3)}$. The MFMC estimator achieves more than two orders of magnitude speedup compared to the single-fidelity Monte Carlo estimator that uses the high-fidelity model alone. The MFMC estimator that combines the high-fidelity model $f^{(1)}$ with two low-fidelity models $f^{(2)}, f^{(3)}$ achieves speedups of more than six orders of magnitude. The large increase in speedup is obtained because the low-fidelity output random variable $f^{(2)}(X)$ and the high-fidelity output random variable $f^{(1)}(X)$ are significantly higher correlated than the low-fidelity output random variable $f^{(3)}(X)$ and $f^{(1)}(X)$, while both low-fidelity models have about the same costs. Table 6 reports the numbers of evaluations of the low- and high-fidelity models for various estimators. The MFMC estimator with models $f^{(1)}, f^{(3)}$ evaluates the high-fidelity model significantly more often than the MFMC estimator with all three models $f^{(1)}, f^{(2)}, f^{(3)}$. Furthermore, by adding model $f^{(2)}$, the MFMC estimator $f^{(1)}, f^{(2)}, f^{(3)}$ seems to make better use of model $f^{(3)}$, because $f^{(3)}$ is more often evaluated in the MFMC estimator with all three models than in the MFMC estimator with $f^{(1)}, f^{(3)}$. This emphasizes that the interaction between the models is what drives the efficiency of MFMC, rather than each of the models separately [32].

V. Conclusions

We demonstrated multifidelity uncertainty propagation with MFMC on two numerical examples. MFMC guarantees unbiased estimators, in contrast to typical techniques that use a low-fidelity model alone. MFMC optimally decides




Fig. 9 Flutter problem with three inputs: Combining the low-fidelity model $f^{(3)}$ and the high-fidelity model $f^{(1)}$ into an MFMC estimator leads to a speedup of more than two orders of magnitude compared to using the high-fidelity model alone. Combining all three models $f^{(1)}, f^{(2)}, f^{(3)}$ leads to an even higher speedup. While $f^{(2)}$ and $f^{(3)}$ have about the same costs, the output random variable corresponding to the low-fidelity model $f^{(2)}$ is higher correlated to the high-fidelity output random variable $f^{(1)}(X)$ than $f^{(3)}(X)$, and therefore the MFMC estimator with all three models achieves a higher speedup than the MFMC estimator that uses $f^{(1)}$ and $f^{(3)}$ only.

Table 6 Flutter problem with three inputs: The table reports the number of times the low- and the high-fidelity models are evaluated in the MFMC estimators. The MFMC estimator that uses $f^{(1)}, f^{(3)}$ performs significantly more high-fidelity model evaluations than the MFMC estimator that combines all three models $f^{(1)}, f^{(2)}, f^{(3)}$.

online costs [s]	Monte Carlo	MFMC with $f^{(1)}, f^{(3)}$		MFMC with $f^{(1)}, f^{(2)}, f^{(3)}$		
	#evals $f^{(1)}$	#evals $f^{(1)}$	#evals $f^{(3)}$	#evals $f^{(1)}$	#evals $f^{(2)}$	#evals $f^{(3)}$
7.56×10^1	5	4.00×10^0	3.46×10^6	3.00×10^0	2.56×10^6	1.70×10^8
1.51×10^2	10	9.00×10^0	6.93×10^6	7.00×10^0	5.13×10^6	3.39×10^8
7.56×10^2	50	4.90×10^1	3.46×10^7	3.50×10^1	2.56×10^7	1.70×10^9
1.51×10^3	100	9.90×10^1	6.93×10^7	7.10×10^1	5.13×10^7	3.39×10^9
7.56×10^3	500	4.97×10^2	3.46×10^8	3.57×10^2	2.56×10^8	1.70×10^{10}

how often to evaluate each of the available models based on estimates of the evaluation costs and the correlation coefficients of the low- and high-fidelity models. Rough estimates of the correlation coefficients obtained from pilot runs are sufficient, because the unbiasedness of the MFMC estimator is independent of the error in the estimated correlation coefficients. MFMC is applicable to any type of low-fidelity model. We derived data-fit surrogate models with procedures that are built in MATLAB and achieved orders of magnitude speedups in our examples. Extensions of the MFMC method to estimate higher-order moments of $f(X)$, e.g., the variance of $f(X)$, are available [40]. A MATLAB implementation of MFMC is available on GitHub <https://github.com/pehersto/mfmc>.

For the flutter study, we used a wing model for which the flutter solution seemed unique and well-behaved. The applicability of the MFMC approach to more complicated aeroelastic models, e.g., of complete aircraft, should be assessed. This assessment would be interesting in situations where there are multiple flutter modes, when establishing the dominant mode may be difficult, owing either to the direct nature of the flutter analysis used here, or to the potential shifting of dominance between fidelity levels. Furthermore, in more realistic flutter studies, the number of uncertain inputs is high, which means that the construction of data-fit surrogate models (e.g., kriging) suffers from the curse of dimensionality. MFMC can be applied to simplified-physics models, e.g., full-potential computational fluid dynamics versus Navier-Stokes solutions, to cope with these high-dimensional problems. A detailed study is necessary to evaluate how effective MFMC is with simplified-physics models, for which typically the speedup compared to the

high-fidelity model is significantly lower than the speedup obtained with data-fit low-fidelity models and response surface surrogates.

Acknowledgments

The first and third author acknowledge support of the AFOSR MURI on multi-information sources of multi-physics systems under Award Number FA9550-15-1-0038 (Dr. Jean-Luc Cambier, Program Officer). The second author acknowledges support of the AFOSR Computational Mathematics Program (Dr. Jean-Luc Cambier, Program Officer).

References

- [1] Forrester, A., Sóbester, A., and Keane, A., *Engineering design via surrogate modelling: a practical guide*, Wiley, 2008.
- [2] Forrester, A. I. J., and Keane, A. J., “Recent advances in surrogate-based optimization,” *Progress in Aerospace Sciences*, Vol. 45, No. 1–3, 2009, pp. 50–79.
- [3] Krige, D. G., “A Statistical Approach to Some Basic Mine Valuation Problems on the Witwatersrand,” *Journal of the Chemical, Metallurgical and Mining Society of South Africa*, Vol. 52, No. 6, 1951, pp. 119–139.
- [4] Cressie, N., “The origins of kriging,” *Mathematical Geology*, Vol. 22, No. 3, 1990, pp. 239–252.
- [5] Rasmussen, C., and Williams, C., *Gaussian Processes for Machine Learning*, MIT Press, 2006.
- [6] Rumpfkeil, M. P., Hanazaki, K., and Beran, P. S., “Construction of Multi-Fidelity Locally Optimized Surrogate Models for Uncertainty Quantification,” *19th AIAA Non-Deterministic Approaches Conference*, AIAA SciTech Forum, American Institute of Aeronautics and Astronautics, 2017.
- [7] Rumpfkeil, M. P., and Beran, P. S., “Construction of Dynamic Multifidelity Locally Optimized Surrogate Models,” *AIAA Journal*, Vol. 55, No. 9, 2017, pp. 3169–3179.
- [8] Sirovich, L., “Turbulence and the dynamics of coherent structures,” *Quarterly of Applied Mathematics*, Vol. 45, 1987, pp. 561–571.
- [9] Rozza, G., Huynh, D., and Patera, A., “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, 2007, pp. 1–47.
- [10] Gugercin, S., Antoulas, A., and Beattie, C., “ H_2 Model Reduction for Large-Scale Linear Dynamical Systems,” *SIAM Journal on Matrix Analysis and Applications*, Vol. 30, No. 2, 2008, pp. 609–638.
- [11] Benner, P., Gugercin, S., and Willcox, K., “A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems,” *SIAM Review*, Vol. 57, No. 4, 2015, pp. 483–531.
- [12] Anttonen, J. S. R., King, P. I., and Beran, P. S., “POD-Based reduced-order models with deforming grids,” *Mathematical and Computer Modelling*, Vol. 38, No. 1, 2003, pp. 41 – 62.
- [13] Lucia, D. J., Beran, P. S., and Silva, W. A., “Reduced-order modeling: new approaches for computational physics,” *Progress in Aerospace Sciences*, Vol. 40, No. 1–2, 2004, pp. 51 – 117.
- [14] Majda, A. J., and Gershgorin, B., “Quantifying uncertainty in climate change science through empirical information theory,” *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 107, No. 34, 2010, pp. 14958–14963.
- [15] Ng, L. W. T., and Willcox, K., “Monte Carlo Information-Reuse Approach to Aircraft Conceptual Design Optimization Under Uncertainty,” *Journal of Aircraft*, Vol. 53, No. 2, 2016, pp. 427–438.
- [16] Peherstorfer, B., Willcox, K., and Gunzburger, M., “Survey of multifidelity methods in uncertainty propagation, inference, and optimization,” *SIAM Review*, 2017. (accepted).
- [17] Eldred, M. S., Ng, L. W. T., Barone, M. F., and Domino, S. P., “Multifidelity Uncertainty Quantification Using Spectral Stochastic Discrepancy Models,” *Handbook of Uncertainty Quantification*, edited by R. Ghanem, D. Higdon, and H. Owhadi, Springer International Publishing, Cham, 2016, pp. 1–45.
- [18] Narayan, A., Gittelson, C., and Xiu, D., “A Stochastic Collocation Algorithm with Multifidelity Models,” *SIAM Journal on Scientific Computing*, Vol. 36, No. 2, 2014, pp. A495–A521.

- [19] Teckentrup, A. L., Jantsch, P., Webster, C. G., and Gunzburger, M., “A Multilevel Stochastic Collocation Method for Partial Differential Equations with Random Input Data,” *SIAM/ASA Journal on Uncertainty Quantification*, Vol. 3, No. 1, 2015, pp. 1046–1074.
- [20] Babuška, I., Nobile, F., and Tempone, R., “A Stochastic Collocation Method for Elliptic Partial Differential Equations with Random Input Data,” *SIAM Journal on Numerical Analysis*, Vol. 45, No. 3, 2007, pp. 1005–1034.
- [21] Nobile, F., Tempone, R., and Webster, C. G., “A Sparse Grid Stochastic Collocation Method for Partial Differential Equations with Random Input Data,” *SIAM Journal on Numerical Analysis*, Vol. 46, No. 5, 2008, pp. 2309–2345.
- [22] Gunzburger, M. D., Webster, C. G., and Zhang, G., “Stochastic finite element methods for partial differential equations with random input data,” *Acta Numerica*, Vol. 23, 2014, pp. 521–650.
- [23] Nelson, B. L., “On control variate estimators,” *Computers & Operations Research*, Vol. 14, No. 3, 1987, pp. 219–225.
- [24] Pasupathy, R., Schmeiser, B. W., Taaffe, M. R., and Wang, J., “Control-variate estimation using estimated control means,” *IIE Transactions*, Vol. 44, No. 5, 2012, pp. 381–385.
- [25] Boyaval, S., “A fast Monte-Carlo method with a reduced basis of control variates applied to uncertainty propagation and Bayesian estimation,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 241–244, 2012, pp. 190–205.
- [26] Boyaval, S., and Lelièvre, T., “A variance reduction method for parametrized stochastic differential equations using the reduced basis paradigm,” *Communications in Mathematical Sciences*, Vol. 8, No. 3, 2010, pp. 735–762.
- [27] Tracey, B., Wolpert, D., and Alonso, J. J., “Using Supervised Learning to Improve Monte-Carlo Integral Estimation,” *AIAA Journal*, Vol. 51, No. 8, 2013, pp. 2015–2023.
- [28] Heinrich, S., “Multilevel Monte Carlo Methods,” *Large-Scale Scientific Computing*, edited by S. Margenov, J. Waśniewski, and P. Yalamov, No. 2179 in Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2001, pp. 58–67.
- [29] Giles, M., “Multi-level Monte Carlo path simulation,” *Operations Research*, Vol. 56, No. 3, 2008, pp. 607–617.
- [30] Vidal-Codina, F., Nguyen, N., Giles, M., and Peraire, J., “A model and variance reduction method for computing statistical outputs of stochastic elliptic partial differential equations,” *Journal of Computational Physics*, Vol. 297, 2015, pp. 700 – 720.
- [31] Ng, L. W. T., and Willcox, K., “Multifidelity approaches for optimization under uncertainty,” *International Journal for Numerical Methods in Engineering*, Vol. 100, No. 10, 2014, pp. 746–772.
- [32] Peherstorfer, B., Willcox, K., and Gunzburger, M., “Optimal model management for multifidelity Monte Carlo estimation,” *SIAM Journal on Scientific Computing*, Vol. 38, No. 5, 2016, pp. A3163–A3194.
- [33] Jasa, J. P., Hwang, J. T., and Martins, J. R. R. A., “Open-source coupled aerostructural optimization using Python,” *Structural and Multidisciplinary Optimization*, 2018. (Submitted).
- [34] Gray, J., Moore, K., and Naylor, B., “OpenMDAO: An Open Source Framework for Multidisciplinary Analysis and Optimization,” *13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, Multidisciplinary Analysis Optimization Conferences, American Institute of Aeronautics and Astronautics, 2010.
- [35] Stanford, B., and Beran, P. S., “Direct flutter and limit cycle computations of highly flexible wings for efficient analysis and optimization,” *Journal of Fluids and Structures*, Vol. 36, No. Supplement C, 2013, pp. 111 – 123.
- [36] Tang, D., and Dowell, E. H., “Experimental and Theoretical Study on Aeroelastic Response of High-Aspect-Ratio Wings,” *AIAA Journal*, Vol. 39, No. 8, 2001, pp. 1430–1441.
- [37] Beran, P. S., Stanford, B. K., and Wang, K. G., “Fast prediction of flutter and flutter sensitivities,” *International Forum on Aeroelasticity and Structural Dynamics*, 2017, pp. 1–20.
- [38] Beran, P., Stanford, B., and Schrock, C., “Uncertainty Quantification in Aeroelasticity,” *Annual Review of Fluid Mechanics*, Vol. 49, No. 1, 2017, pp. 361–386.
- [39] Tran, C., and Petot, D., “Semi-empirical model for the dynamic stall of airfoils in view to the application to the calculation of responses of a helicopter blade in forward flight,” *Vertica*, Vol. 5, 1981, pp. 35–53.
- [40] Qian, E., Peherstorfer, B., O’Malley, D., Vesselinov, V. V., and Willcox, K., “Multifidelity Monte Carlo estimation of variance and sensitivity indices,” Tech. Rep. ACDL Technical Report TR-2017-2, Massachusetts Institute of Technology, 2017.