

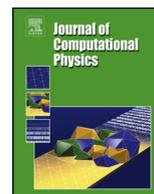


ELSEVIER

Contents lists available at ScienceDirect

## Journal of Computational Physics

www.elsevier.com/locate/jcp



# Practical error bounds for a non-intrusive bi-fidelity approach to parametric/stochastic model reduction



Jerrad Hampton<sup>a</sup>, Hillary R. Fairbanks<sup>b</sup>, Akil Narayan<sup>c</sup>, Alireza Doostan<sup>a,\*</sup>

<sup>a</sup> Smead Aerospace Engineering Sciences Department, University of Colorado, Boulder, CO 80309, USA

<sup>b</sup> Department of Applied Mathematics, University of Colorado, Boulder, CO 80309, USA

<sup>c</sup> Department of Mathematics, and Scientific Computing and Imaging (SCI) Institute, University of Utah, Salt Lake City, UT 84112, USA

## ARTICLE INFO

### Article history:

Received 12 September 2017

Received in revised form 6 March 2018

Accepted 6 April 2018

Available online 10 April 2018

### Keywords:

Matrix interpolative decomposition

Uncertainty quantification

Low-rank approximation

Multi-fidelity approximation

Bi-fidelity approximation

Parametric model reduction

## ABSTRACT

For practical model-based demands, such as design space exploration and uncertainty quantification (UQ), a high-fidelity model that produces accurate outputs often has high computational cost, while a low-fidelity model with less accurate outputs has low computational cost. It is often possible to construct a bi-fidelity model having accuracy comparable with the high-fidelity model and computational cost comparable with the low-fidelity model. This work presents the construction and analysis of a non-intrusive (i.e., sample-based) bi-fidelity model that relies on the low-rank structure of the map between model parameters/uncertain inputs and the solution of interest, if exists. Specifically, we derive a novel, pragmatic estimate for the error committed by this bi-fidelity model. We show that this error bound can be used to determine if a given pair of low- and high-fidelity models will lead to an accurate bi-fidelity approximation. The cost of this error bound is relatively small and depends on the solution rank. The value of this error estimate is demonstrated using two example problems in the context of UQ, involving linear and non-linear partial differential equations.

Published by Elsevier Inc.

## 1. Introduction

In many practical contexts, an ideal computational method accurately recovers the physical phenomena that it is tasked to model, and does so in a computationally efficient manner that allows repeated calculations of quantities of interest (QoI's) for a large number of different input parameters. When this is possible, the QoI's may be understood across a range of input parameters, or similarly for certain distributions of input parameters. The latter case is connected with the field of parametric uncertainty quantification (UQ), and is fundamental for understanding the propagation of uncertainty through complex models often via the construction of the map between the inputs and QoI's. A growing area of interest in model-based simulations is a fast construction of such a mapping, which also forms our motivation here. For physical systems that are expensive to simulate, such a task may not be straightforward and may provide a significant challenge as it often requires a tradeoff between computational cost and accuracy.

In practice, the model which accurately represents the underlying physics – the *high-fidelity* model – may be expensive, and many repeated evaluations are likely computationally infeasible. Alternatively, one may consider the use of a *low-fidelity* model: This is a model with lower accuracy which is cheaper to evaluate – relative to the high-fidelity model – but not

\* Corresponding author.

E-mail address: alireza.doostan@colorado.edu (A. Doostan).

necessarily as accurate. The topic of this paper is a *bi-fidelity* approach that leverages both models and inherits the strengths of each.

In this work, we assume the existence of a low- and a high-fidelity model that each identify the same QoI in a discrete sense, i.e., as a vector, possibly with different sizes and levels of accuracy. For example, a fine spatial discretization of a heat equation, while costly, may lead to accurate temperature predictions, but a coarse spatial discretization gives inaccurate estimates that are generated more quickly. The notion of accuracy here is contextual, both with respect to the metric and the desired quantitative tolerance. Mathematically, we take the high-fidelity model as one that gives an approximation to the *truth* to within a desired accuracy and ideally with smallest possible computational cost. It is critical that we assume such a high-fidelity model is still expensive to evaluate, so that methods that would require many simulations from this model are impractical. We also assume that it is (much) cheaper to sample from the low-fidelity model so that compiling a larger number of samples from the low-fidelity model is possible. The goal of a (non-intrusive) bi-fidelity, or more generally multi-fidelity, model is to use relatively few samples from the high-fidelity model and a larger number of samples from the low-fidelity model to generate approximations to the QoI – at arbitrary input samples – with accuracies comparable to that of the high-fidelity model.

Early work in multi-fidelity methods for parametric problems has a decidedly statistical flavor. Such work emerged in the field of geosciences with the application of Gaussian process regression – also known as kriging or co-kriging in the multivariate case – in a multi-fidelity setting. In [1], an autoregressive scheme is considered to generate a Gaussian process approximation of the output of the most expensive model from nested observations of multiple, less expensive models. In this autoregressive scheme, a given model is approximated by the lower-fidelity model through a multiplicative constant and an additive Gaussian process correction term. In a related work, [2] creates a Gaussian process surrogate to the low-fidelity model, which is subsequently adjusted via samples of the high-fidelity model. The adjustment is done by training a similar autoregressive model as in [1], consisting in a linear – in input parameters – mapping of the low-fidelity surrogate and an additive discrepancy term modeled as a Gaussian process. The resulting high-fidelity surrogate is also a Gaussian process. Extensions of multi-fidelity Gaussian process modeling for applications in optimization are explored in [3,4], and recursive methods for improved accuracy or computational complexity are developed in [5–10], among other work.

There has also been significant work over the last several years on reformulating other prominent UQ methods in a multi-fidelity (or multi-level) framework. These methods generally rely on a relatively large number of low-fidelity samples along with a (much) smaller number of high-fidelity samples to build an additive and/or multiplicative correction of the low-fidelity model. The adjusted low-fidelity model will then serve as a multi-fidelity approximation to the high-fidelity model. In particular, in multi-fidelity polynomial chaos expansions (PCEs) [11–14], the high-fidelity solution is estimated by first determining the PCE of the low-fidelity solution, and subsequently the PCE of a correction term is generated from a small set of low-fidelity and high-fidelity samples. In multi-level Monte Carlo (MC) [15–17], a sequence of fine-to-coarse level (or grid) models are considered and the expectation of the QoI is estimated via a telescoping summation. These methods apply MC to the coarsest level QoI, and then add the MC estimates of the QoI differences, a.k.a. corrections, over each pair of consecutive grids. Assuming the corrections over finer grids become smaller, fewer realizations of the finer (high-fidelity) grid QoI are needed to compute the correction expectations, thus resulting in a reduced overall cost.

With an aim of constructing a reduced polynomial representation of the high-fidelity solution, the work in [18,19] builds a small size polynomial basis – in the random inputs – identified by a Karhunen-Loève expansion of the low-fidelity solution. The high-fidelity solution is then approximated via Galerkin projection of the governing equations on the span of the reduced polynomial basis – as opposed to the standard PCE basis – resulting in a much smaller size system of equations to solve. The key assumption in this bi-fidelity construction is that the largest eigenvalues of the covariance matrix of the discrete solution decay quickly; that is, the covariance matrix and, therefore, the solution are low-rank. A non-intrusive implementation of this work is found in [20].

Of most relevance to the approach of this study, [21,22] considers the low-rank approximation of the discretized QoI  $\mathbf{v} \in \mathbb{R}^M$ ,

$$\mathbf{v}(\boldsymbol{\mu}) \approx \sum_{\ell=1}^r \mathbf{v}(\boldsymbol{\mu}^{(\ell)}) c_{\ell}(\boldsymbol{\mu}). \quad (1)$$

The approximation above is of the same type constructed in so-called *reduced basis* representations in the (parametric) model reduction literature [23–27]. However, the procedure considered in [21,22] and here is disparate from the algorithmic strategies in the reduced basis community. In (1),  $\mathbf{v}$  at an arbitrary parametric input  $\boldsymbol{\mu} \in \mathbb{R}^d$  is represented in a basis consisting of  $r \ll M$  realizations of  $\mathbf{v}$  – corresponding to selected input samples  $\boldsymbol{\mu}^{(\ell)}$  – via the coefficients  $c_{\ell}(\boldsymbol{\mu})$ . The work in [21,22] relies on a greedy procedure applied to low-fidelity realizations of  $\mathbf{v}$  to identify the set of  $\boldsymbol{\mu}^{(\ell)}$ . The coefficients  $c_{\ell}(\boldsymbol{\mu})$  are computed via an interpolatory condition on the low-fidelity model. The same interpolation rule, i.e., the same coefficients  $c_{\ell}(\boldsymbol{\mu})$ , is then used with the high-fidelity counterpart of the reduced basis snapshots and this results in the bi-fidelity construction (1). In this manner, the high-fidelity model is evaluated only at a (small) number of input samples that is equal to the size  $r$  of the reduced basis. Unlike in standard reduced basis techniques this approach does not perform Galerkin (or Petrov–Galerkin) projections of the high-fidelity equations, thus allowing the use of legacy codes in a black-box fashion.

With a different algorithm to construct the reduced basis and the interpolation rule, the work in [28] and [29] examines a similar bi-fidelity approximation on problems involving heat transfer in a ribbed channel and prediction of pressure coefficient over a family of NACA airfoils, respectively. The strategy in [28] was used in [30] to build control variates within a multi-level MC framework.

### 1.1. Contribution of this work

While in practice the above multi-fidelity methods have shown promising results, there is a lack of convenient tools to verify the convergence of the multi-fidelity solution and understand the role of key factors affecting the convergence. These tools must also enable a practitioner to determine, *a priori* and with relatively small cost, whether or not a given low-fidelity model will result in accurate multi-fidelity approximations. To this end, our primary contribution is the derivation of an error bound that is appropriate for methods as in [21,22,28,30,29,31] for bi-fidelity approximation, with a particular emphasis on the role of the rank of the  $\mathbf{v}$  ensemble in the approximation error. This bound is derived specifically for the bi-fidelity procedure in [28,30,29,31] and relies on quantities that are easily estimated with a number of high-fidelity samples that scales favorably with the rank of the  $\mathbf{v}$  ensemble. As a result, the associated error estimate is efficient, easy to compute, and reliably conservative, making it desirable in practical contexts. This method is also presented in its full generality while applied to two examples that demonstrate the potential for practical utility of this error estimate.

The remainder of this manuscript is organized into three main sections. Section 2 reviews the mathematical notation and procedure for the bi-fidelity approximation. Section 2.4 presents our main result, which focuses on bounding the bi-fidelity approximation error. To show the accuracy of this bound, Section 3 provides two numerical experiments and illustrates the utility of the error estimate. Finally, Section 4 gives a short summary of the conclusions of this study. In Appendix A, we provide the proof of the main error bound introduced in Section 2.4.

## 2. Method detail

As in [21,22,28,30,29,31], we consider a bi-fidelity construction of a QoI admitting a low-rank representation of the form (1). We presume here that these QoI's are functions defined over the spatial or temporal domain of the problem or can be computed from such functions. While the QoI's could contain the solution over the entire domain of the problem, e.g., temperature over the physical domain, this is often unnecessary and a reduction to some summary of the solution, e.g., temperature along a boundary, can explain the behavior of interest. In a broad sense, this low-rank heuristic implies the existence of a limited number of behaviors for  $\mathbf{v}$  to exhibit, and that for any particular set of inputs  $\boldsymbol{\mu}$ , the corresponding  $\mathbf{v}$  is nearly a linear reconstruction of  $\mathbf{v}$  exhibiting these different traits. In this way, even  $\mathbf{v}$  for possibly highly non-linear problems may be approximated, with some tolerable error, by a linear combination of relatively few basis vectors. This low-rank assumption is actually more than a heuristic: It can be codified via the mathematical notion of  $n$ -widths [32], and QoI ensembles for differential equations can indeed exhibit this low-rank property [33].

This formulation raises two primary concerns: How do we identify the basis vectors, and how do we identify the appropriate linear combinations of these basis vectors? For computational feasibility, we require solutions to both these concerns that require neither a large number of high-fidelity realizations of  $\mathbf{v}$ , nor modification to the simulation codes.

To present our approach for addressing these concerns, and without loss of generality, we consider the approximation of a collection of  $N$  high-fidelity realizations of the QoI, instead of individual realizations as in (1). We denote the  $k$ th realization of the QoI obtained via the high- and low-fidelity models by  $\mathbf{v}_H^{(k)} \in \mathbb{R}^M$  and  $\mathbf{v}_L^{(k)} \in \mathbb{R}^m$ , respectively. These vectors may be of different sizes, but depend on the same parametric input sample  $\boldsymbol{\mu}^{(k)}$ . For the examples presented in Section 3, we consider  $\boldsymbol{\mu}$  as a random variable, and  $\mathbf{v}_H^{(k)}$  and  $\mathbf{v}_L^{(k)}$  are drawn in an MC fashion from the joint probability distribution on  $\boldsymbol{\mu}$ . When no probabilistic information on  $\boldsymbol{\mu}$  is available,  $\mathbf{v}_H^{(k)}$  and  $\mathbf{v}_L^{(k)}$  may be generated, e.g., from uniformly distributed samples of  $\boldsymbol{\mu}$ , over some prescribed region of permissible  $\boldsymbol{\mu}$ .

Our analysis relies on arranging  $\{\mathbf{v}_L^{(k)}\}$  and  $\{\mathbf{v}_H^{(k)}\}$  into matrices, denoted  $\mathbf{H} \in \mathbb{R}^{M \times N}$  and  $\mathbf{L} \in \mathbb{R}^{m \times N}$ , to refer to high-fidelity and low-fidelity data, respectively, where the corresponding columns of  $\mathbf{H}$  and  $\mathbf{L}$  are the QoI realizations at the same input samples. That is,  $\mathbf{H}$  and  $\mathbf{L}$  have the same number of columns,  $N$ , but generally have different numbers of rows  $M$  and  $m$ , respectively. The low-rank assumption on the  $\mathbf{v}_H$  ensemble implies that  $\mathbf{H}$  admits a low-rank factorization. We assume the same for  $\mathbf{L}$  but do not require the same rank for  $\mathbf{H}$  and  $\mathbf{L}$ .

**Remark 1.** For the interest of presentation and analysis we refer to the full high-fidelity matrix  $\mathbf{H}$ . However, in practice, we require having access to only  $r \ll N$  ( $r$  being the rank of  $\mathbf{v}_H$  as in (1)) high-fidelity realizations to generate this bi-fidelity approximation; see Section 2.2. To evaluate the bound on the resulting approximation error, we require a number of high-fidelity realizations that is slightly larger than  $r$ , as illustrated in the examples of Section 3.

### 2.1. Low-rank factorization of low-fidelity data

The method used to identify a reduced basis and an interpolation rule for  $\mathbf{H}$  (or  $\mathbf{v}_H$ ) relies on identifying the same for the low-fidelity data  $\mathbf{L}$ . For this purpose, we form the rank  $r$  *interpolative decomposition* of  $\mathbf{L}$  [34–36] using the rank-revealing

QR factorization [37]. The revealed rank of this factorization is our reduced rank  $r$ , and this provides the rank- $r$  reduced decomposition

$$\mathbf{LP} \approx \mathbf{Q}_L \mathbf{R}_L, \tag{2}$$

for an  $N \times N$  permutation matrix  $\mathbf{P}$ . Then  $\mathbf{Q}_L$  has  $r$  orthonormal columns, and  $\mathbf{R}_L$  has  $r$  rows and is upper triangular. A column partitioning of  $\mathbf{R}_L$  into an  $r \times r$  upper triangular matrix  $\mathbf{R}_L^{(11)}$  and an  $r \times (N - r)$  matrix  $\mathbf{R}_L^{(12)}$ , and replacing  $\mathbf{R}_L^{(12)}$  with the approximation  $\mathbf{R}_L^{(12)} \approx \mathbf{R}_L^{(11)} \mathbf{Z}$  for some  $\mathbf{Z}$  leads to

$$\mathbf{LP} \approx \mathbf{Q}_L \mathbf{R}_L^{(11)} [\mathbf{I} \mid \mathbf{Z}], \tag{3}$$

where  $\mathbf{I}$  is the  $r \times r$  identity matrix. When  $\mathbf{R}_L^{(11)}$  is ill-conditioned, [34] suggests a solution  $\mathbf{Z}$  with minimum  $\|\mathbf{Z}\|_F$ , where  $\|\cdot\|_F$  denotes the Frobenius norm. The rank  $r$  factorization (3) can be rewritten as

$$\begin{aligned} \mathbf{L} &\approx \mathbf{L}(r) [\mathbf{I} \mid \mathbf{Z}] \mathbf{P}^T; \\ &= \mathbf{L}(r) \mathbf{C}_L; \end{aligned} \tag{4}$$

$$:= \hat{\mathbf{L}}, \tag{5}$$

where  $\mathbf{L}(r) = \mathbf{Q}_L \mathbf{R}_L^{(11)}$  – referred to as the *column skeleton* of  $\mathbf{L}$  [34] – contains the first  $r$  columns of  $\mathbf{LP}$  and  $\mathbf{C}_L := [\mathbf{I} \mid \mathbf{Z}] \mathbf{P}^T$ . In words, (4) identifies a set of  $r$  columns of  $\mathbf{L}$ , i.e., a reduced basis for  $\mathbf{L}$ , along with a coefficient matrix  $\mathbf{C}_L$  specifying the linear combination of the basis vectors to approximate all columns in  $\mathbf{L}$ . Hence the name *interpolative decomposition*. If we identify the entries of  $\mathbf{C}_L$  with the coefficients  $c_\ell(\boldsymbol{\mu})$  in (1), then (4) is the matrix formulation of (1) for  $N$  low-fidelity realizations  $\{\mathbf{v}_L^{(k)}\}$ .

For the sake of completeness, we next report some properties of the factorization (4) from [35], which we will use later. For detailed analyses of the interpolative decomposition, the interested reader is referred to [34–36].

**Lemma 1.** (Lemma 3.1 of [35].) Let  $\|\cdot\|$  denote the matrix induced  $\ell_2$  norm. For any positive integer  $r \leq \min\{m, N\}$ ,

1.  $\|\mathbf{C}_L\| \leq \sqrt{r(N - r) + 1}$ ,
2.  $\mathbf{L} = \hat{\mathbf{L}}$ , when  $r = m$  or  $r = N$ , and
3.  $\|\mathbf{L} - \hat{\mathbf{L}}\| \leq \sqrt{r(N - r) + 1} \sigma_{r+1}$  when  $r < \min\{m, N\}$ , where  $\sigma_{r+1}$  is the  $(r + 1)$ st greatest singular value of  $\mathbf{L}$ .

**Remark 2.** In practice, the approximation rank  $r$  is not known *a priori* and thus the pivoted Gram–Schmidt step involved in the QR factorization (2) is continued with larger ranks until  $\|\mathbf{L} - \hat{\mathbf{L}}\|$  reaches a specified accuracy.

**Remark 3.** Although not investigated here, other factorizations such as those in [38–41] are equally applicable to produce an interpolative decomposition of  $\mathbf{L}$ .

### 2.2. High-fidelity approximation via basis update and low-fidelity interpolation

In addition to a rank  $r$  factorization of  $\mathbf{L}$ , the permutation matrix  $\mathbf{P}$  in the interpolative decomposition (4) identifies the set of  $r$  input samples  $\boldsymbol{\mu}^{(k)}$  corresponding to the basis vectors in  $\mathbf{L}(r)$ . This identification is one primary reason for employing an interpolative decomposition instead of the more widely-used and  $\ell_2$ -optimal singular value decomposition (SVD) of  $\mathbf{L}$ . As in [21,22,28,30,29,31], our approach to approximate  $\mathbf{H}$  is to replace  $\mathbf{L}(r)$  with  $\mathbf{H}(r)$ , a matrix of vectors from the high-fidelity model corresponding to the same input samples  $\boldsymbol{\mu}^{(k)}$  and with the same arrangement as in  $\mathbf{L}(r)$ . Stated differently, we consider the high-fidelity counterpart of  $\mathbf{L}(r)$  as the reduced basis for  $\mathbf{H}$  (or  $\mathbf{v}_H^{(k)}$ ). In this case we set the bi-fidelity, rank  $r$  approximation of  $\mathbf{H}$ , denoted by  $\hat{\mathbf{H}}$ , to

$$\hat{\mathbf{H}} := \mathbf{H}(r) \mathbf{C}_L, \tag{6}$$

where  $\mathbf{C}_L$  is the coefficient matrix computed from the low-fidelity approximation (4). In other words, the interpolation rule learned from the low-fidelity model is applied to approximate the high-fidelity realizations. In the case that we have a construction of this form, where the difference between  $\hat{\mathbf{L}}$  and  $\hat{\mathbf{H}}$  is limited to the changes only in the basis vectors associated with the low- and high- fidelity models, we say that  $\hat{\mathbf{L}}$  is *lifted* to the approximation  $\hat{\mathbf{H}}$ . For this paper we say that  $\hat{\mathbf{L}}$  and  $\hat{\mathbf{H}}$  are *corresponding estimates* (of  $\mathbf{L}$  and  $\mathbf{H}$ , respectively). Our analysis presented next applies to any such corresponding estimates.

**Remark 4** (Utility of  $\hat{\mathbf{H}}$  for UQ). While outside the scope of the present study, the bi-fidelity realizations  $\hat{\mathbf{v}}_H^{(k)}$ , i.e., the columns of  $\hat{\mathbf{H}}$ , may be used in place of  $\mathbf{v}_H^{(k)}$  in methods such as MC simulation or its variants, sparse grid stochastic collocation [42], non-intrusive PCE [43–46], etc., to approximate the statistics of  $\mathbf{v}_H$  or perform sensitivity analysis. The accuracy of such estimates, however, depends on the accuracy of the bi-fidelity estimates  $\hat{\mathbf{v}}_H^{(k)}$ .

### 2.3. Number of low- and high-fidelity samples

Notice that the bi-fidelity construction (6) relies on  $r$  high-fidelity samples. Therefore, when  $\mathbf{v}_H$  admits a low-rank representation, i.e.,  $r$  is small, only a small number of high-fidelity samples are needed. To generate a rank- $r$  approximation of  $\mathbf{L}$ , a theoretical lower bound on the number of required low-fidelity model evaluations  $N$  (and  $m$ ) is  $r$ . In practice, however, a considerably larger number of low-fidelity samples may be needed to identify a representative reduced basis. Finding an optimal  $N$  to reduce the cost of low-fidelity model evaluations depends on the problem at hand, and on how  $\mathbf{v}_L$  varies with respect to  $\boldsymbol{\mu}$ . For scenarios where the low-fidelity model is not substantially cheaper than its high-fidelity counterpart, adaptive sampling strategies – similar to those in [47,48] – can be employed to reduce the number of low-fidelity model evaluations. This is, however, outside the scope of this manuscript.

### 2.4. Convergence analysis and its practical application

We are now prepared to state the main result that bounds the bi-fidelity error,  $\|\hat{\mathbf{H}} - \mathbf{H}\|$ , where  $\|\cdot\|$  is the matrix induced  $\ell_2$  norm. However, we first present a high-level discussion on our analysis approach. The core of our convergence results rely on the assumption that there exists a matrix  $\mathbf{T}$ , with  $\|\mathbf{T}\|$  bounded, so that  $\mathbf{H} \approx \mathbf{T}\mathbf{L}$ . The difference in these matrices is denoted  $\mathbf{E}$ , and we require that  $\|\mathbf{E}\|$  is small. Restated, our bound involves the matrix relation given by

$$\mathbf{H} = \mathbf{T}\mathbf{L} + \mathbf{E}. \tag{7}$$

The matrix  $\mathbf{T}$  is essentially a lifting operator from low-fidelity discretization space to high-fidelity discretization space. Our theory outlined in Appendix A computes this matrix explicitly, but does not require any knowledge about the discretizations or assumptions between the low- and high-fidelity models. We emphasize that  $\mathbf{T}$  is only needed for theoretical analysis and our error bounds do not require its construction.

It is straightforward to verify that  $\|\hat{\mathbf{H}} - \mathbf{H}\|$  can be bounded as

$$\begin{aligned} \|\mathbf{H} - \hat{\mathbf{H}}\| &\leq \|\mathbf{H} - \mathbf{T}\mathbf{L}\| + \|\mathbf{T}\mathbf{L} - \mathbf{T}\hat{\mathbf{L}}\| + \|\mathbf{T}\hat{\mathbf{L}} - \hat{\mathbf{H}}\|; \\ &\leq \|\mathbf{H} - \mathbf{T}\mathbf{L}\| + \|\mathbf{T}\| \|\mathbf{L} - \hat{\mathbf{L}}\| + \|\mathbf{T}\mathbf{L} - \mathbf{H}\| \|\mathbf{C}_L\|; \\ &\leq (1 + \|\mathbf{C}_L\|) \|\mathbf{E}\| + \|\mathbf{T}\| \|\mathbf{L} - \hat{\mathbf{L}}\|. \end{aligned} \tag{8}$$

When  $\|\mathbf{T}\|$  is bounded and  $\|\mathbf{E}\|$  is small, (8) suggests an accurate bi-fidelity estimate as  $\|\mathbf{C}_L\|$  is bounded, and  $\|\mathbf{L} - \hat{\mathbf{L}}\|$  is small given that  $\mathbf{L}$  is low-rank; see Lemma 1. In general, there is no guarantee of the existence of a mapping  $\mathbf{T}$  of  $\mathbf{L}$  to  $\mathbf{H}$  that gives an  $\mathbf{E}$  with small norm. We therefore seek to establish a condition on  $\mathbf{L}$  (in relation to  $\mathbf{H}$ ) to ensure the existence of such  $\mathbf{T}$  and  $\mathbf{E}$  matrices.

For a finite  $\tau \geq 0$ , define

$$\epsilon(\tau) := \lambda_{\max}(\mathbf{H}^T \mathbf{H} - \tau \mathbf{L}^T \mathbf{L}), \tag{9}$$

where  $\lambda_{\max}(\cdot)$  denotes the largest eigenvalue of a matrix, and  $\mathbf{H}^T \mathbf{H}$  and  $\mathbf{L}^T \mathbf{L}$  are the Gramians of the high- and low-fidelity matrices, respectively. We describe in Appendix A that if  $\epsilon(\tau)$  is small enough then matrices  $\mathbf{T}$  and  $\mathbf{E}$  with our desired properties can be constructed.

For a fixed  $\tau$ ,  $\epsilon(\tau)$  has various interpretations. It is the smallest  $\epsilon$  such that

$$\epsilon \mathbf{I} + \tau \mathbf{L}^T \mathbf{L} - \mathbf{H}^T \mathbf{H} \tag{10}$$

is a positive semi-definite matrix. Second,  $\epsilon(\tau)$  has a geometrical interpretation that demonstrates its significance in this context of bi-fidelity approximation. The above equation implies that  $\epsilon(\tau)$  is the smallest value such that for any  $\mathbf{x} \in \mathbb{R}^N$ ,

$$\|\mathbf{H}\mathbf{x}\|^2 \leq \tau \|\mathbf{L}\mathbf{x}\|^2 + \epsilon(\tau) \|\mathbf{x}\|^2.$$

This guarantees that for every  $\mathbf{x}$ ,  $\sqrt{\tau} \mathbf{L}\mathbf{x}$  is further from the origin than  $\mathbf{H}\mathbf{x}$ , to within a margin governed by  $\epsilon(\tau)$  and  $\|\mathbf{x}\|$ . This ensures that every realized  $\mathbf{H}\mathbf{x}$  can be reached by rotating  $\mathbf{L}\mathbf{x}$  through the application of a scaling related to  $\tau$ , and adding a small correction which is bounded based on  $\epsilon(\tau)$  and  $\|\mathbf{x}\|$ . The relationship between these interpretations is further discussed in Appendix A, and we note this rotation, scaling, and additive correction is essentially utilized to prove the following theorem completing the bound in (8).

**Theorem 1.** For any  $\tau \geq 0$ , let  $\epsilon(\tau)$  be as in (9). Let  $\hat{\mathbf{H}}$  and  $\hat{\mathbf{L}}$  be corresponding estimates of rank  $r$  with coefficients  $\mathbf{C}_L$ , and let  $\sigma_k$  denote the  $k$ th largest singular value of  $\mathbf{L}$ . Then,

$$\|\mathbf{H} - \hat{\mathbf{H}}\| \leq \min_{\tau, k \leq \text{rank}(\mathbf{L})} \rho_k(\tau) \tag{11a}$$

$$\rho_k(\tau) := \left[ (1 + \|\mathbf{C}_L\|) \sqrt{\tau \sigma_{k+1}^2 + \epsilon(\tau)} + \|\mathbf{L} - \hat{\mathbf{L}}\| \sqrt{\tau + \epsilon(\tau) \sigma_k^{-2}} \right]. \tag{11b}$$

When  $k = \text{rank}(\mathbf{L})$ , we set  $\sigma_{k+1} = 0$ .

The proof of this theorem is given in Appendix A; here we provide insight and implications of its result.

An actual evaluation of the bound requires estimating  $\tau$  and  $\epsilon(\tau)$ , which can be achieved by minimizing the bound over a range of  $k$  and  $\tau$ ; see Algorithm 1. A naïve attempt to estimate  $\epsilon(\tau)$  from (9) seems to require the entire matrix  $\mathbf{H}$ . However, we propose in Section 2.5 a procedure that uses only a relatively small number of high- and low-fidelity realizations. Notice that aside from  $\epsilon(\tau)$ , the bound (11) has no dependence on the high-fidelity model. Let

$$B_1 = (1 + \|\mathbf{C}_L\|) \sqrt{\tau \sigma_{k+1}^2 + \epsilon(\tau)}; \quad (12)$$

$$B_2 = \|\mathbf{L} - \hat{\mathbf{L}}\| \sqrt{\tau + \epsilon(\tau) \sigma_k^{-2}}. \quad (13)$$

If the approximation is accurate, the expressions in (12) and (13) will be small. Specifically, observe that  $B_1$  is small when both  $\tau \sigma_{k+1}^2 \ll 1$  and the optimal  $\epsilon(\tau)$  is small. Recall  $\|\mathbf{C}_L\| \leq \sqrt{r(N-r)+1}$  from Lemma 1. When, for a given  $k$ ,  $\sigma_{k+1}$  is small, i.e., the singular values of  $\mathbf{L}$  decay quickly, the magnitude of  $B_1$  is governed by the optimal  $\epsilon(\tau)$ . To see that  $B_2$  is small, note that if the low-rank reconstruction is accurate for the low-fidelity model, then  $\|\mathbf{L} - \hat{\mathbf{L}}\| \leq \sqrt{r(N-r)+1} \sigma_{r+1}$  is small following Lemma 1. Given this,  $B_2$  is also small when  $\sigma_{k+1}/\sigma_k \ll 1$ ,  $\tau \sigma_{k+1}^2 \ll 1$ , and the optimal  $\epsilon(\tau)$  is small. We note that the requirements  $\sigma_{k+1}/\sigma_k \ll 1$  and  $\tau \sigma_{k+1}^2 \ll 1$  highlight the significance of the low-rank assumption on  $\mathbf{L}$  in the success of this bi-fidelity construction. Specifically, a low-rank approximation will have a  $k$  such that  $\sigma_{k+1}$  is significantly smaller than  $\sigma_k$ .

Using Lemma 1, the bound (11) can be simplified to

$$\|\mathbf{H} - \hat{\mathbf{H}}\| \leq \min_{\tau, k \leq \text{rank}(\mathbf{L})} \left[ \sqrt{r(N-r)+1} \left( 1 + \frac{\sigma_{k+1}}{\sigma_k} \right) \sqrt{\tau \sigma_k^2 + \epsilon(\tau)} \right].$$

In our numerical experiments, we use the sharper error estimate in (11).

We now address the computation of an optimal  $\epsilon(\tau)$ , or, more precisely an optimal pair  $(\tau, \epsilon(\tau))$ . Specifically we consider an estimate  $\hat{\epsilon}(\tau)$  with relatively few evaluations of the high-fidelity model. Such an estimate allows us to approximate the bound (11) on  $\|\mathbf{H} - \hat{\mathbf{H}}\|$ , which measures the quality of the bi-fidelity approximation.

### 2.5. Estimating $\epsilon(\tau)$ using limited high-fidelity data

The definition of  $\epsilon(\tau)$  in (9) depends on full realizations of the high-fidelity model which is untenable in many practical situations and its direct computation would significantly detract from the utility of the bi-fidelity error estimate. Consider Gramian matrices and an estimate  $\hat{\epsilon}$  defined, for a normalizing constant  $c$ , by

$$\mathbf{G}_H := \mathbf{H}^T \mathbf{H}; \quad (14)$$

$$\mathbf{G}_L := \mathbf{L}^T \mathbf{L}; \quad (15)$$

$$\hat{\epsilon}(\tau) := c \lambda_{\max}(\hat{\mathbf{G}}_H - \tau \hat{\mathbf{G}}_L), \quad (16)$$

for estimates  $\hat{\mathbf{G}}_H$  and  $\hat{\mathbf{G}}_L$  of  $\mathbf{G}_H$  and  $\mathbf{G}_L$ , respectively. We note the similarity of (9) and (16), where the latter makes use of quantities with the hat notation. Of key practical note is that  $\mathbf{G}_H$  inherits the low rank nature of  $\mathbf{H}$ , so that the estimate of  $\epsilon(\tau)$  can be accurate with a small number of samples. Recall that the number of columns of  $\mathbf{H}$  and  $\mathbf{L}$  is  $N$ . Let  $n$  denote a small number of columns that are sub-sampled from  $\mathbf{H}$  and  $\mathbf{L}$ , and  $\hat{\mathbf{G}}_H$  and  $\hat{\mathbf{G}}_L$  are the Gramian matrices associated with these  $n$  columns. We set  $c = N/n$  as the normalizing constant throughout the remainder of this work. With these Gramian estimates, we can construct  $\hat{\epsilon}(\tau)$  for any  $\tau$  using (16). To approximate the bound in (11), we may replace  $\epsilon(\tau)$  with  $\hat{\epsilon}(\tau)$ , and then identify the minimum achieved value over a range of  $k$  and  $\tau$  values, noting that the search over  $k$  requires only additionally knowing the appropriate singular values of  $\mathbf{L}$ . This requires  $n$  high-fidelity samples (to form the Gramian matrix  $\hat{\mathbf{G}}_H$ ). Our numerical results in Section 3 show empirically that  $n$  slightly larger than  $r$  is sufficient to estimate the optimal pair  $(\tau, \epsilon(\tau))$ . Notice, however, that the low-fidelity error  $\|\mathbf{L} - \hat{\mathbf{L}}\|$  in (11) is calculated using all  $N$  low-fidelity samples.

Algorithm 1 summarizes our proposed approach to evaluate the bound (11) as an estimate for the bi-fidelity error  $\|\mathbf{H} - \hat{\mathbf{H}}\|$ .

---

#### Algorithm 1 Algorithm for estimating the error bound (11).

---

- 1: Initialize a vector of values for  $\tau \geq 0$ .
  - 2: Use  $n$  high- and low-fidelity realizations to set the Gramian matrices  $\hat{\mathbf{G}}_L$  and  $\hat{\mathbf{G}}_H$  as explained in Section 2.5.
  - 3: For all values of  $\tau$  evaluate the corresponding  $\hat{\epsilon}(\tau)$  using (16).
  - 4: For each  $k \in \{0, \dots, \text{rank}(\mathbf{L})\}$  evaluate (11) for all pairs  $(\tau, \hat{\epsilon}(\tau))$  computed in Step 3 above.
  - 5: Choose the minimum value achieved by (11) over all  $k$  and  $(\tau, \hat{\epsilon}(\tau))$  as an estimate for  $\|\mathbf{H} - \hat{\mathbf{H}}\|$ .
-

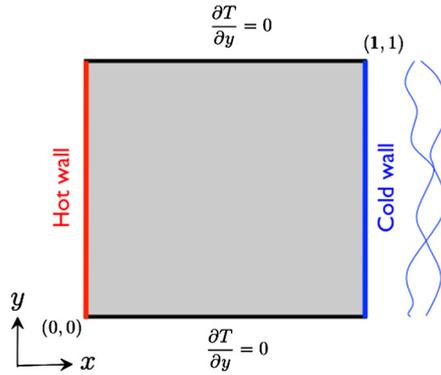


Fig. 1. Schematic of the heat driven cavity flow problem, reproduced from Fig. 5 of [30].

### 3. Numerical examples

To investigate various aspects of the proposed bi-fidelity approximation and the associated error estimate, we consider two practically motivated problems.

#### 3.1. Test case 1: heat driven cavity flow

A practical case for consideration comes from temperature-driven fluid flow in a cavity [49–54,30,55], where the QoI is the steady state heat flux along the hot wall; see Fig. 1. The left vertical wall has a random temperature  $T_h$ , referred to as the hot wall, while the right vertical wall, referred to as the cold wall, has a spatially varying stochastic temperature  $T_c < T_h$  with constant mean  $\bar{T}_c$ . Both horizontal walls are adiabatic. The reference temperature and the reference temperature difference are defined as  $T_{ref} = \bar{T}_c$  and  $\Delta T_{ref} = T_h - \bar{T}_c$ , respectively. Under small temperature difference assumption, i.e., the Boussinesq approximation, the normalized governing equations are given by, [49],

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla p + \frac{\text{Pr}}{\sqrt{\text{Ra}}} \nabla^2 \mathbf{u} + \text{Pr} \Theta \mathbf{e}_y, \\ \nabla \cdot \mathbf{u} &= 0, \\ \frac{\partial \Theta}{\partial t} + \nabla \cdot (\mathbf{u} \Theta) &= \frac{1}{\sqrt{\text{Ra}}} \nabla^2 \Theta, \end{aligned} \tag{17}$$

where  $\mathbf{e}_y$  is the unit vector  $(0, 1)$ ,  $\mathbf{u} = (u, v)$  is velocity vector field,  $\Theta = (T - T_{ref})/\Delta T_{ref}$  is normalized temperature,  $p$  is pressure, and  $t$  is time. Zero velocity boundary conditions on all walls (in both directions) are assumed. For more details on the normalization of the variables in (17), we refer the interested reader to [51,49]. Prandtl and Rayleigh numbers are defined, respectively, as  $\text{Pr} = \nu/\alpha$  and  $\text{Ra} = g\tau\Delta T_{ref}L^3/(\nu\alpha)$ . Specifically,  $L$  is the length of the cavity,  $g$  is gravitational acceleration,  $\nu$  is kinematic viscosity,  $\alpha$  is thermal diffusivity, and the coefficient of thermal expansion is given by  $\tau$ . In this example, we set  $g = 10$ ,  $L = 1$ ,  $\tau = 0.5$ , and  $\text{Pr} = 0.71$ . We use a finite volume method for the discretization of (17).

##### 3.1.1. Sources of uncertainty

On the cold wall, a temperature distribution with stochastic fluctuations is applied as

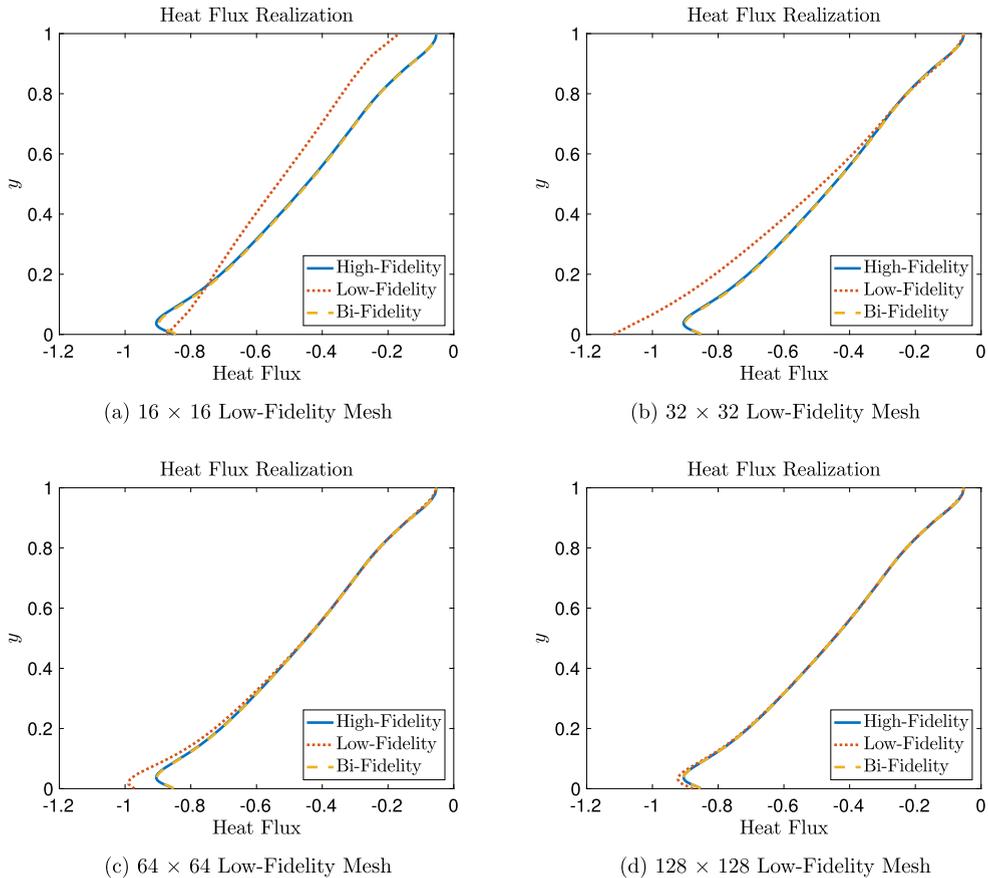
$$T(x = 1, y) = \bar{T}_c + T'(y), \tag{18}$$

where  $\bar{T}_c = 100$  is a constant expected temperature. The fluctuation  $T'(y)$  is given by the truncated Karhunen-Loève-like expansion

$$T'(y) = \sigma_T \sum_{i=1}^{d_T} \sqrt{\lambda_i} \phi_i(y) \mu_i, \tag{19}$$

where  $d_T = 50$  and  $\sigma_T = 2$ . Here, each  $\mu_i$  is assumed to be an independent and identically distributed uniform random variable on  $[-1, 1]$ , and  $\{\lambda_i\}_{i=1}^{d_T}$  and  $\{\phi_i(y)\}_{i=1}^{d_T}$  are the  $d_T$  largest eigenvalues and the corresponding eigenfunctions of the exponential covariance kernel

$$C_{TT}(y_1, y_2) = \exp\left(-\frac{|y_1 - y_2|}{l_T}\right), \tag{20}$$



**Fig. 2.** Realizations of heat flux along the hot wall for a randomly selected input  $\mu$ . Shown are the low-fidelity, high-fidelity, and rank  $r = 10$  bi-fidelity estimates for the various low-fidelity models.

with correlation length  $l_T = 0.15$ . An example of the cold boundary condition is shown in Fig. 1. The temperature on the hot wall  $T_h$  is also assumed to be random and uniformly distributed over  $[105, 109]$ . Finally, we consider the viscosity  $\nu$  to be uniformly distributed over  $[0.004, 0.01]$ . In total, the dimension of the random input  $\mu = (\mu_1, \dots, \mu_{50}, T_h, \nu)$  is 52. The QoI, heat flux along the hot wall, is represented by a vector of heat flux values at the grid points (along the hot wall) corresponding the high- or low-fidelity meshes.

### 3.1.2. Bi-fidelity approximation and error bound estimate

To compute realizations of the QoI, we use finite volume discretizations of various resolutions to solve (17), where the highest resolution, a  $256 \times 256$  grid point mesh, is used as the high-fidelity model, while relatively coarse meshes of  $128 \times 128$ ,  $64 \times 64$ ,  $32 \times 32$ , and  $16 \times 16$  grid points are used for the corresponding low-fidelity model. All meshes are spatially uniform. We consider bi-fidelity approximations of various ranks  $r$ , corresponding to the number of basis vectors in the column skeleton matrices  $\mathbf{H}(r)$  and  $\mathbf{L}(r)$  introduced in Section 2.4. We also consider bounds derived from computations of optimal pair  $(\tau, \hat{\epsilon}(\tau))$  as explained in Section 2.5. These estimates are drawn from  $n$  randomly generated high-fidelity samples, i.e.,  $n$  randomly selected columns of  $\mathbf{H}$ , and we consider estimates computed from various sample sizes  $n$ .

Fig. 2 displays realizations of heat flux along the hot wall for a random sample of  $\mu$ , obtained by the high-fidelity, low-fidelity, and rank  $r = 10$  bi-fidelity models. We observe the close agreement between the bi- and high-fidelity solutions even when the  $16 \times 16$  mesh is used as the low-fidelity model. Fig. 3 provides four histograms of the low- and bi-fidelity error to gauge the performance of all 100 QoI realizations. For (a)–(d), which use increasingly refined meshes for the low-fidelity model, we see an improvement of the bi-fidelity performance over that of the associated low-fidelity model.

We next consider the error bound efficacy of using  $\hat{\epsilon}(\tau)$  from (16), in place of  $\epsilon(\tau)$  from (9), in Theorem 1. This efficacy is defined as the ratio of the error estimate over the actual error in the approximation of  $\mathbf{H}$  so that an efficacy of 1 implies the error estimate is exact, and an efficacy greater than 1 implies that the computed error estimate is in fact a bound. Fig. 4 shows this error efficacy for various low-fidelity models, approximation ranks  $r$ , and sample sizes  $n$ . For each pair  $(n, r)$  the reported error ratio is the average of 30 error ratios, each computed from independent sets of  $n$  high-fidelity samples. We note that the error efficacy is greater than one whenever the sample size  $n$  is greater than the approximation rank  $r$ , and that, as long as this condition holds, the error efficacy is less than ten, implying that the error bound is off by less than

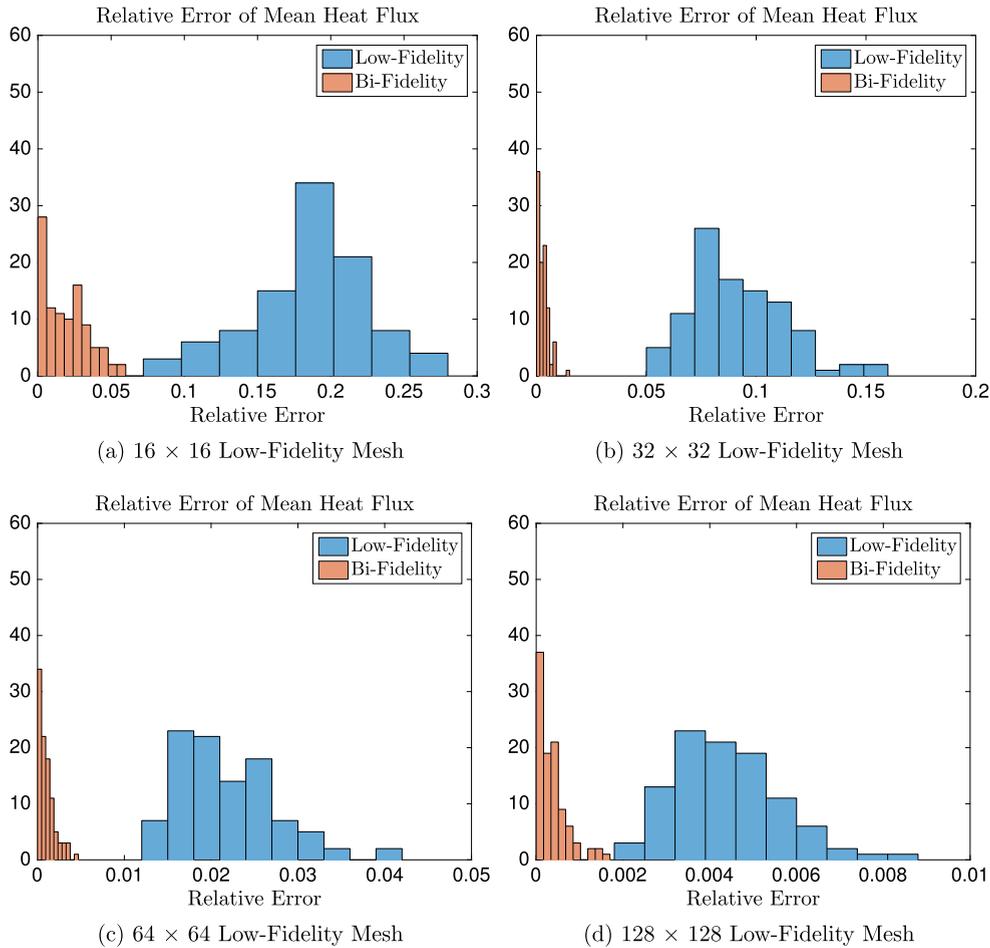


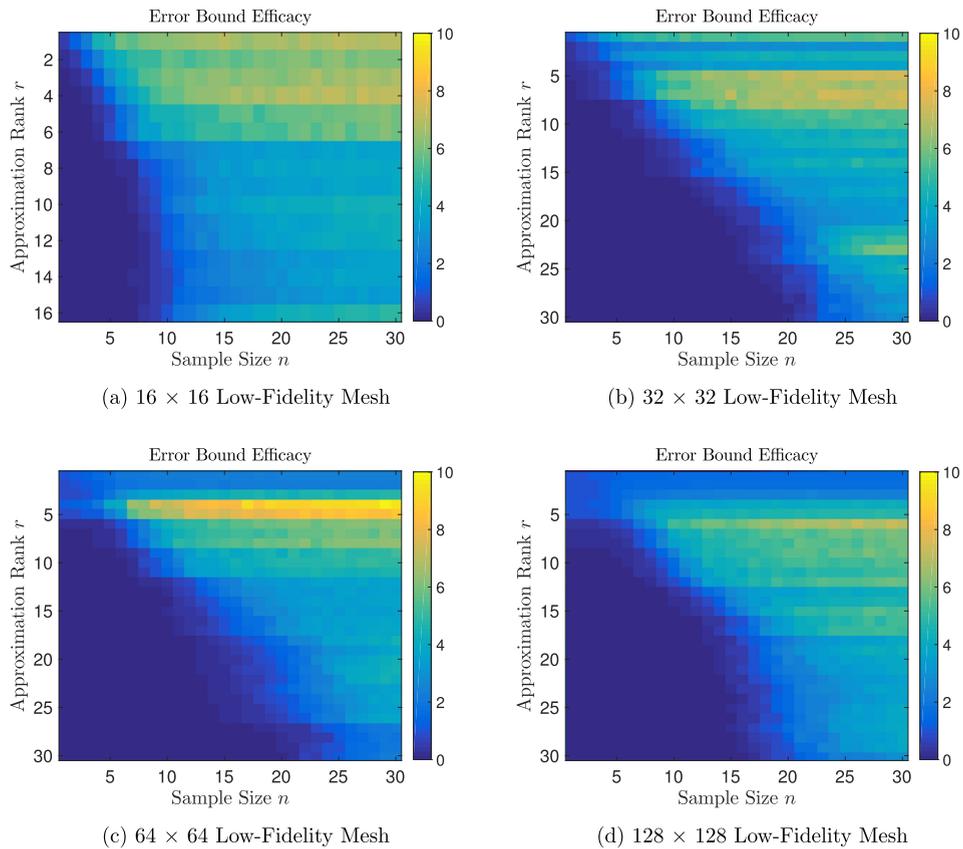
Fig. 3. Histograms of low-fidelity and rank  $r = 10$  bi-fidelity errors normalized by  $\|\mathbf{H}\|$  for various low-fidelity models from 100 data samples.

one order of magnitude. This is of great utility in determining either a sample size  $n$  for a given approximation rank  $r$  or, conversely, determining an admissible rank of approximation for a given sample size for which the computed error estimate is still valid. These estimates indicate that the error bounds are generally conservative with relatively few samples.

To address the estimation of (11) numerically, we consider in Fig. 5 how  $\hat{\epsilon}(\tau)$  in (16) behaves as a function of  $\tau$  for four combinations of rank  $r$  and sample size  $n$ . We mark points that optimize (11) with the corresponding bi-fidelity rank  $r$ , and note the reduction in the optimal values of  $\tau$  and  $\hat{\epsilon}(\tau)$  as the mesh of the low-fidelity model is refined. Fig. 5 (a) and (b), which provide values of  $\hat{\epsilon}(\tau)$  for  $r = 5$ , show that while estimates of  $\hat{\epsilon}(\tau)$  differ for the two values of  $n$ , points of  $(\tau, \hat{\epsilon}(\tau))$  that optimize (11) are close. In contrast, Fig. 5 (c) and (d), which provide the same set of results but for  $r = 10$ , show dissimilar optimal points. These differences correspond to the ability to describe points of  $(\tau, \epsilon(\tau))$ . Recall, from the results of Fig. 4 that small values of  $n$  can result in underestimating the bound. For the scenario with  $r = 10$  and  $n = 10$ , these error bound estimates are not reliable, which is further highlighted by optimal values occurring at  $\hat{\epsilon}(\tau) = 0$  in Fig. 5 (c).

We can also explore the behavior of  $\rho_k(\tau)$  as a function of  $\tau$  in (11b), noting that minimizing over  $\tau$  and  $k$  yields our main estimate (11a). This function is seen for different low-fidelity model meshes in Fig. 6. We observe that as the low-fidelity meshes are refined,  $\tau$  converges to 1, corresponding to a broadly defined convergence of low-fidelity models to the high-fidelity model. We also note, that there is not a high level of sensitivity with regards to  $\tau$ , and that a large range of  $\tau$  produces comparable bounds. Fig. 6 also illustrates the benefit that the optimization of  $\tau$  need not be particularly accurate, and that the general shape of these curves is smooth enough that derivative-based optimization can be employed, limiting the number of  $\tau$  for which  $\hat{\epsilon}(\tau)$  needs to be evaluated.

It is worth highlighting that the heuristic employed here is that the size of the matrices used in  $\hat{\epsilon}(\tau)$  need only depend on the approximation rank  $r$  and scale nearly as such. As a result, the cost of evaluating  $\hat{\epsilon}(\tau)$  will typically not be prohibitively large, as illustrated in Fig. 7 by the rapid convergence of the optimal  $\hat{\epsilon}(\tau)$  with respect to  $n$ . Similar to Fig. 4, this figure also shows that when the sample size is too small (relative to  $r$ ), e.g.,  $n = 8$ , the error estimate is below the actual error. Finally, Fig. 8 shows the convergence of the  $r = 10$  bi-fidelity solution as a function of the low-fidelity mesh size. The error bound estimate – relative to  $\mathbf{H}$  – is displayed for two sample sizes of  $n$ , and the low-fidelity error – also



**Fig. 4.** Identification of error bound efficacy, i.e., the (average) ratio between the error estimated from (11) and the true error, for various low-fidelity meshes, approximation ranks  $r$ , and sample sizes  $n$ . For each pair  $(n, r)$  the reported error ratio is the average of 30 ratios, each computed from an independent set of  $n$  high- and low-fidelity samples via Algorithm 1. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

**Table 1**

Description of the parameters of the composite cantilever beam model. The center of the holes are at  $x = \{5, 15, 25, 35, 45\}$ .

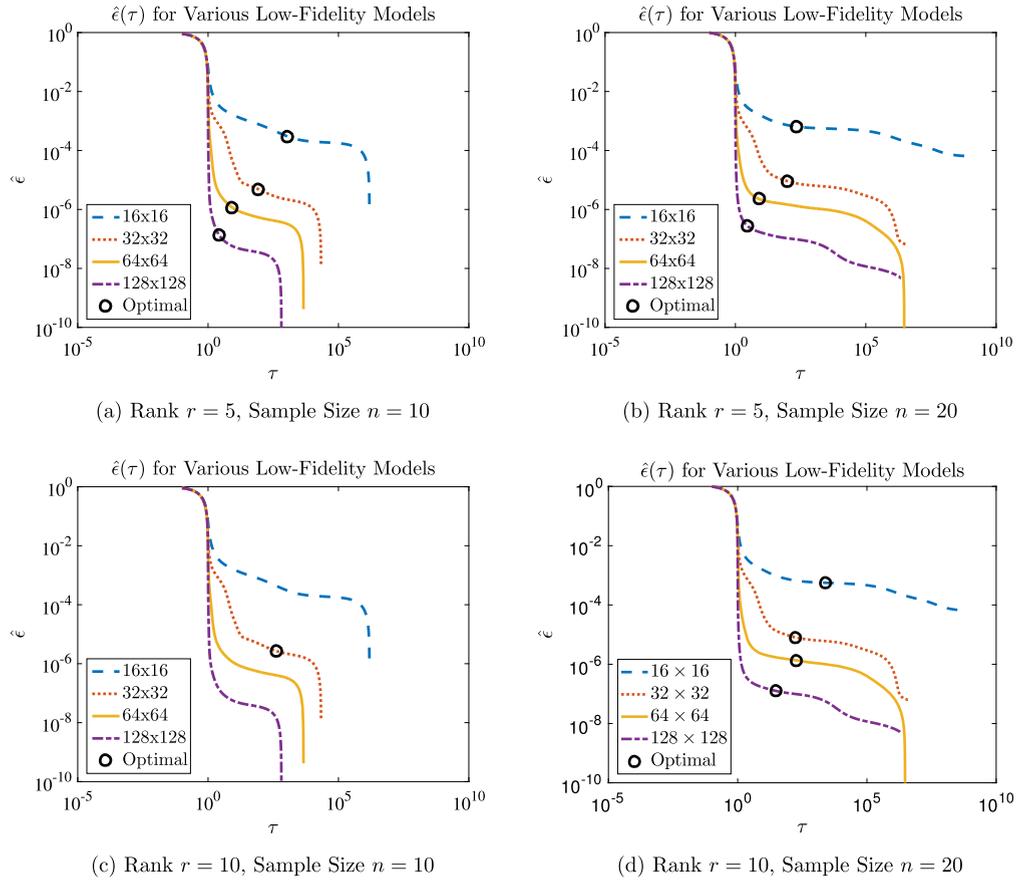
$L$	$h_1$	$h_2$	$h_3$	$w$	$r$	$q$	$E_1$	$E_2$	$E_3$
50	0.1	0.1	5	1	1.5	$U[9, 11]$	$U[0.9e6, 1.1e6]$	$U[0.9e6, 1.1e6]$	$U[0.9e4, 1.1e4]$

relative to  $\mathbf{H}$  – is provided for comparison. Notice that, aside for a single point from the coarsest low-fidelity model, all the estimated errors are below the corresponding low-fidelity errors. In addition, the error bound estimate decays with the bi-fidelity error, as the mesh is refined.

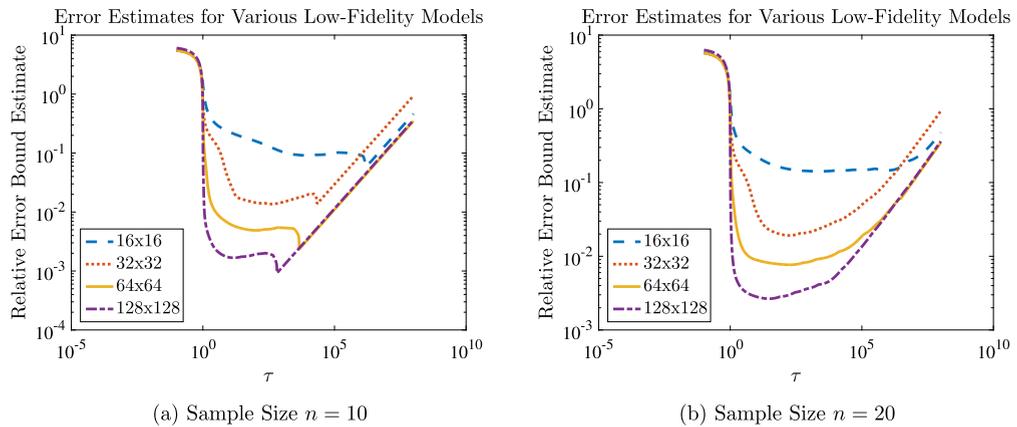
### 3.2. Test case 2: composite beam

For the second test case, we consider the deformation of a plane-stress, cantilever beam with composite cross section and hollow web, as shown in Fig. 9. The Young's moduli of the three components of the cross section as well as the intensity of the applied distributed force on the beam are assumed to be uncertain, and are modeled by independent uniform random variables. Table 1 provides a summary of the parameters of this model along with the description of the uncertain inputs. Here, the QoI is the vertical displacement of the top cord and, in particular, its maximum occurring at the free end. To construct the bi-fidelity approximation, we use realizations of the vertical displacement over the entire top cord given by the low- and high-fidelity models.

Unlike in the previous test case where the low-fidelity model simulated the same physical problem but on a coarser mesh, here the low-fidelity model solves a simpler physical model. Specifically, the high-fidelity model is based on a finite element (FE) discretization of the beam using a triangular mesh with linear elements; see Fig. 10. The mesh is fine enough to describe the web's hollow geometry and achieve mesh independent predictions of the vertical displacement at the free end. The FE simulations are performed using FEniCS [56]. The low-fidelity model is based on Euler-Bernoulli beam theory, where vertical cross sections are assumed to remain planes throughout the deformation and, thus, the shear deformation of the web is ignored. Additionally, we simplify the geometry of the beam by ignoring the circular holes. These simplifications,



**Fig. 5.**  $\hat{\epsilon}(\tau)$  from (16) as a function of  $\tau$  for different low-fidelity model meshes with four combinations of bi-fidelity ranks  $r = 5, 10$  and sample sizes  $n = 10, 20$ . Values which optimize (11) for a rank  $r$  bi-fidelity approximation are marked by black circles. Curves without optimal value markers displayed identify optimal bounds at  $\hat{\epsilon}(\tau) = 0$ .

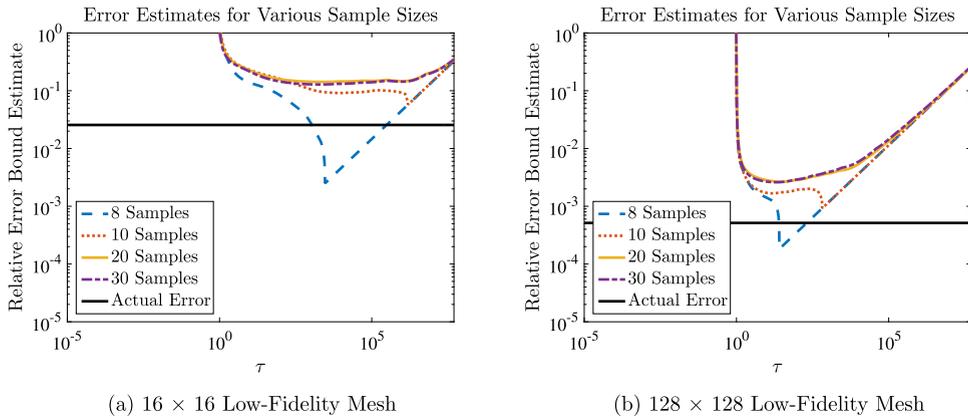


**Fig. 6.** Minimum (over  $k$ ) of error function  $\rho_k(\tau)$  in (11b) for various low-fidelity models using an approximation rank  $r = 10$ . In (a), minimizations occur (from coarsest to finest mesh) at  $k = 13, 8, 9$ , and  $85$ . In (b), minimizations occur (from coarsest to finest mesh) at  $k = 7, 8, 9$ , and  $9$ . Errors are normalized by  $\|\mathbf{H}\|$ .

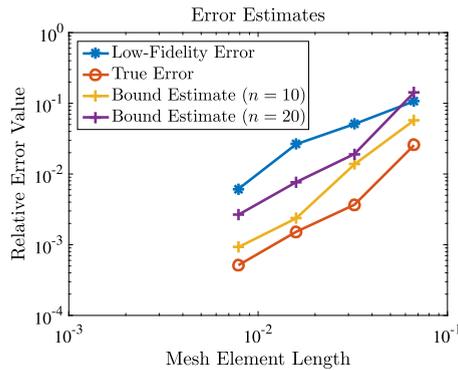
while making the calculations straightforward, lead to *inaccurate* low-fidelity predictions of the displacement (Fig. 11), as the web experiences considerable shear deformation due to the presence of the circular holes.

Following the Euler–Bernoulli theorem, the vertical displacement of the beam  $u(x)$  is given by

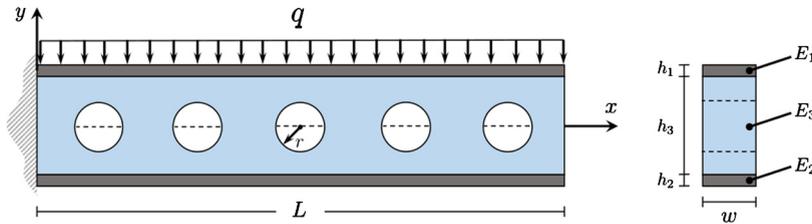
$$EI \frac{d^4 u(x)}{dx^4} = -q, \tag{21}$$



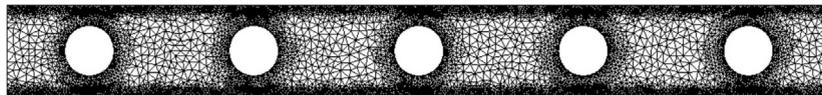
**Fig. 7.** Minimum (over  $k$ ) of error function  $\rho_k(\tau)$  in (11b) for different low-fidelity models using different sample sizes  $n$ . All bi-fidelity solutions are of rank  $r = 10$ . In (a), minimizations occur (in ascending order of  $n$ ) at  $k = 13, 13, 7$ , and  $7$ . In (b), minimizations occur (in ascending order of  $n$ ) at  $k = 85, 85, 9$ , and  $9$ . Errors are normalized by  $\|\mathbf{H}\|$ .



**Fig. 8.** Estimate of (11) for various low-fidelity models and using different sample sizes  $n$ . Errors are normalized by  $\|\mathbf{H}\|$ .



**Fig. 9.** Schematic of the cantilever beam (left) and its composite cross section (right).



**Fig. 10.** Finite element mesh used for high-fidelity simulation of the vertical displacement.

where  $E$  and  $I$  are, respectively, the Young’s modulus and the moment of inertia of an *equivalent* cross section (about its centroid axis) consisting of a single material. Specifically, we let  $E = E_3$  and consider an equivalent cross section in which the width of the top and bottom sections are set to  $w_1 = (E_1/E_3)w$  and  $w_2 = (E_2/E_3)w$ , respectively, while all other dimensions remain as before; see Fig. 9 and Table 1. The solution to (21) for the considered beam is then given by

$$u(x) = -\frac{qL^4}{24EI} \left( \left(\frac{x}{L}\right)^4 - 4\left(\frac{x}{L}\right)^3 + 6\left(\frac{x}{L}\right)^2 \right). \tag{22}$$

We note that the form of the function  $u(x)$  depends only weakly on the uncertain parameters; in particular the parameters only influence the magnitude of  $u(x)$  via the multiplicative coefficient  $qL^4/24EI$ . This implies that the low-fidelity

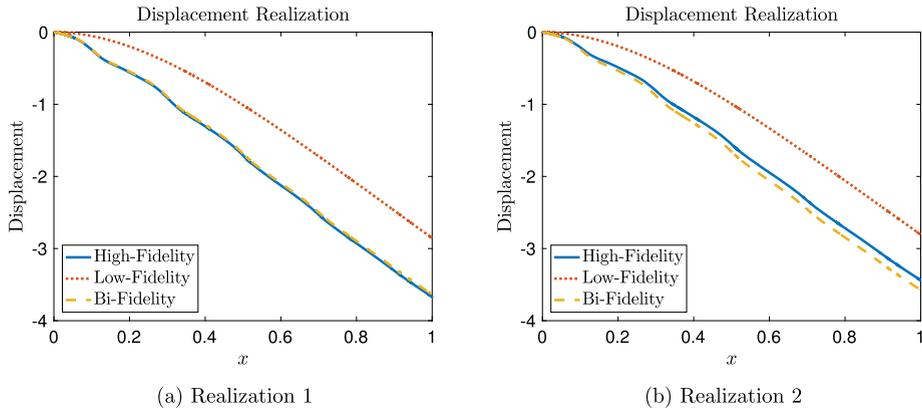


Fig. 11. Realizations of vertical displacement for two randomly selected input parameters  $\mu$ . Shown are the low-fidelity, high-fidelity, and rank  $r = 1$  bi-fidelity estimates.

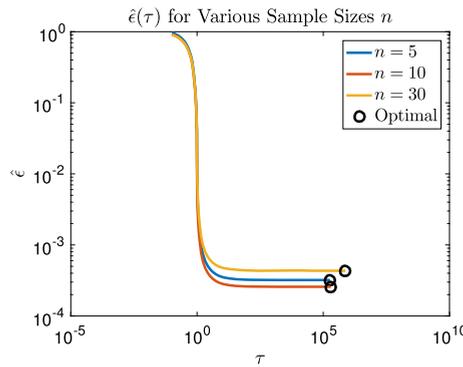


Fig. 12.  $\hat{\epsilon}(\tau)$  from (16) as a function of  $\tau$  for different sample sizes  $n$ . Values which optimize (11) for a rank  $r = 1$  approximation are marked.

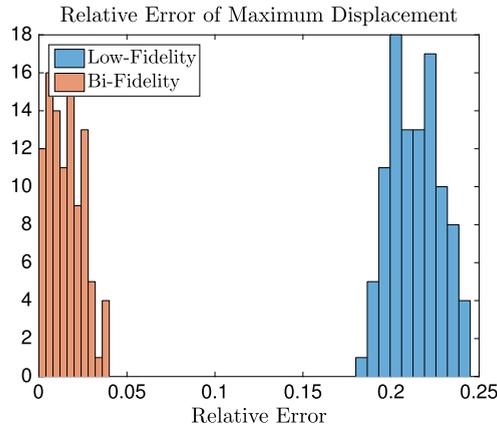
solution is theoretically of rank  $r = 1$  and that only one high-fidelity realization of  $u(x)$  is needed to generate the bi-fidelity approximation. Stated differently, the bi-fidelity approximation cannot be improved by increasing the rank beyond  $r = 1$ . This has significant implications for  $\epsilon(\tau)$  in (9) and the corresponding interpretation of Theorem 1. Specifically, with regards to (11) we note that the only valid choice of  $r$  is  $r = 1$ ; therefore, both  $\sigma_{r+1}$  and  $\|\mathbf{L} - \hat{\mathbf{L}}\|$  are precisely zero so that

$$\|\mathbf{H} - \hat{\mathbf{H}}\| \leq \min_{\tau} (1 + \|\mathbf{C}_L\|) \sqrt{\epsilon(\tau)}.$$

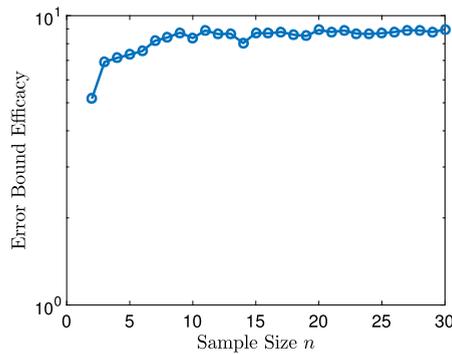
As  $\epsilon$  is a decreasing function of  $\tau$ , we note that the minimizing value of  $\tau$  occurs in the limit as  $\tau \rightarrow \infty$ . This can be numerically observed from Fig. 12. Even though the low-fidelity model is of a minimal rank, the high-fidelity model is itself reasonably well approximated by a representation of rank  $r = 1$ . This is seen in Fig. 11 for two randomly selected choices of input parameters, illustrating a considerably closer agreement between the bi- and high-fidelity samples relative to the low- and high-fidelity samples. Fig. 13 shows histograms of the relative errors of the maximum displacement predictions given by the low- and bi-fidelity models. The histograms are based on 100 random realizations of the maximum displacement. The average error achieved by the bi-fidelity model is roughly an order of magnitude smaller, illustrating that the bi-fidelity approximation is indeed considerably more accurate, in spite of using only one high-fidelity sample (as  $r = 1$ ). In Fig. 14 we see that the error bound stabilizes by around 7 samples and overestimates the true error by a factor 9, which again verifies that a small number of high-fidelity model samples can reliably generate a conservative bound on the error. With regards to the efficacy of the bound (11) and its dependency on  $\tau$  as a function of the sample size  $n$  used to determine an optimal  $\hat{\epsilon}(\tau)$ , we see in Fig. 15 that the bound estimate, and specifically its dependency on  $\tau$ , is robust to  $n$ , converging quickly with respect to  $n$ .

#### 4. Conclusions

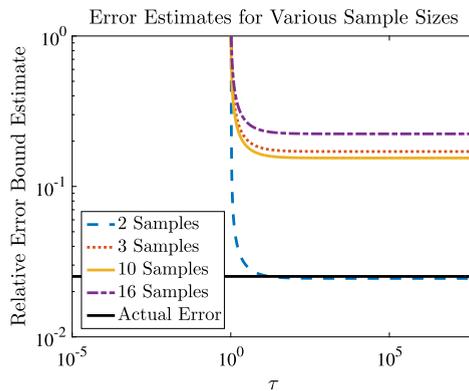
This work is concerned with the error analysis of a bi-fidelity, low-rank approximation technique for a reduced order solution of problems with stochastic or parametric inputs. The bi-fidelity method relies on the solution to a cheaper, lower-fidelity model of the intended expensive, high-fidelity model of the problem, in order to identify a reduced basis and an



**Fig. 13.** Histogram count for the relative error in the maximum displacement for low- and bi-fidelity models from 100 random realizations of the displacement.



**Fig. 14.** Identification of error bound efficacy, i.e., the (average) ratio between the error estimated from (11) and the true error, for various sample sizes  $n$ . For each  $n$  the reported error ratio is the average of 30 ratios, each computed from an independent set of  $n$  high- and low-fidelity samples via Algorithm 1.



**Fig. 15.** Minimum (over  $k$ ) of error function  $\rho_k(\tau)$  in (11b) using different sample sizes  $n$ . Minimization occurs at  $k = 1$  for all four estimates. Errors are normalized by  $\|\mathbf{H}\|$ .

interpolation rule for approximating the high-fidelity solution in its inputs. A novel and practical error bound is presented for this bi-fidelity approximation that is shown to not require many high-fidelity samples beyond the rank of the approximation. This makes the bound an effective tool for practical application of this approximation when a limited computational budget is available. In particular, the error bound allows a practitioner to determine if a given pair of low- and high-fidelity models will lead to an accurate bi-fidelity approximation before investing too many computational resources. The bi-fidelity estimate and the associated error bound are by construction non-intrusive, i.e., sample-based, and thus can use legacy solvers in a black box fashion.

Some aspects of the bi-fidelity approximation, and specifically the efficacy of the derived error bound, are demonstrated on two numerical examples from fluid (non-linear) and solid (linear) mechanics. The numerical results suggest that the proposed error bound achieves a conservative estimate of the true error, while using a number of high-fidelity samples that is slightly larger than the rank of approximation. A precise estimate of the minimum number of high-fidelity samples needed to generate a conservative error bound, however, remains an open problem and requires further research.

When a low-fidelity model is not considerably cheaper than its high-fidelity counterpart, adaptive strategies for sampling the low-fidelity model are needed to identify the reduced basis and the interpolation rule with a minimal number of low-fidelity model evaluations. Additionally, deriving sharper error bounds and extending this work to include data from multiple low-fidelity models are interesting future research directions.

**Acknowledgements**

The work of JH and AN was supported by the DARPA EQuiPS project, award number N660011524053. The work of HF was supported by the United States Department of Energy under the Predictive Science Academic Alliance Program 2 (PSAAP2) at Stanford University. This material is based upon work of AD supported by the U.S. Department of Energy Office of Science, Office of Advances Scientific Computing Research, under Award Number DE-SC0006402, and NSF grant CMMI-145460. AN was partially supported by AFOSR FA9550-15-1-0467.

**Appendix A. Theoretical exposition**

This section is devoted to proving our main theoretical result, Theorem 1. Recall that  $\mathbf{H}$  represents the QoI realization matrix for the high-fidelity model, and  $\mathbf{L}$  the corresponding matrix for the low-fidelity model. Our analysis assumes that these matrices are related through matrices  $\mathbf{T}$  and  $\mathbf{E}$  according to (7). Following the discussions of Section 2.4 and, in particular, the bound in (8), we seek to identify conditions implying that  $\|\mathbf{T}\|$  is bounded and  $\|\mathbf{E}\|$  is small.

For  $\tau \geq 0$ , define

$$\epsilon(\tau) := \arg \min_{\epsilon} \left\{ \forall \mathbf{x} \in \mathbb{R}^N : \|\mathbf{H}\mathbf{x}\|^2 \leq \tau \|\mathbf{L}\mathbf{x}\|^2 + \epsilon \|\mathbf{x}\|^2 \right\}. \tag{A.1}$$

Note that  $\epsilon(\tau)$  is well-defined, is a non-increasing function of  $\tau$  and, satisfies  $\epsilon(\tau) \leq \|\mathbf{H}\|^2$ . Defining  $\epsilon$  and  $\tau$  this way is equivalent to the definition (9) used in Theorem 1. We present this equivalence here as a lemma.

**Lemma 2.** For any  $\tau$ , for  $\epsilon(\tau)$  defined as in (A.1), it follows that

$$\epsilon(\tau) = \lambda_{\max}(\mathbf{H}^T \mathbf{H} - \tau \mathbf{L}^T \mathbf{L}), \tag{A.2}$$

i.e.,  $\epsilon(\tau)$  is the smallest  $\epsilon$  such that  $\tau \mathbf{L}^T \mathbf{L} + \epsilon \mathbf{I} - \mathbf{H}^T \mathbf{H}$  is a semi-positive-definite (SPD) matrix.

**Proof.** For a given  $\tau$ , and  $\epsilon(\tau)$  as from (A.1), it follows that for all nontrivial  $\mathbf{x}$ ,

$$\begin{aligned} \|\mathbf{H}\mathbf{x}\|^2 &\leq \tau \|\mathbf{L}\mathbf{x}\|^2 + \epsilon(\tau) \|\mathbf{x}\|^2; \\ \frac{\tau \|\mathbf{L}\mathbf{x}\|^2 + \epsilon(\tau) \|\mathbf{x}\|^2 - \|\mathbf{H}\mathbf{x}\|^2}{\|\mathbf{x}\|^2} &\geq 0. \end{aligned}$$

By the definition in (A.1),  $\epsilon(\tau)$  is the smallest such number making this inequality hold for all  $\mathbf{x}$ , meaning that the inequality is equality for some  $\mathbf{x}$ , implying

$$\begin{aligned} \lambda_{\min}(\tau \mathbf{L}^T \mathbf{L} + \epsilon(\tau) \mathbf{I} - \mathbf{H}^T \mathbf{H}) &= 0; \\ \lambda_{\min}(\tau \mathbf{L}^T \mathbf{L} - \mathbf{H}^T \mathbf{H}) &= -\epsilon(\tau); \\ \lambda_{\max}(\mathbf{H}^T \mathbf{H} - \tau \mathbf{L}^T \mathbf{L}) &= \epsilon(\tau), \end{aligned}$$

completing the proof.  $\square$

We now identify the specific matrices  $\mathbf{T}$  which we consider for our analysis of (7).

**Lemma 3.** Let the SVD of  $\mathbf{L}$  be given by

$$\mathbf{L} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T.$$

Let  $\mathcal{V}_k$  be the linear subspace spanned by the first  $k$  singular vectors in  $\mathbf{v}$ , i.e., the right singular vectors associated with the  $k$  largest singular values, and let  $\mathbf{V}_k$  be the matrix of those singular vectors. Let

$$\begin{aligned} \mathbf{P}_{\mathcal{V}_k} &= \mathbf{V}_k \mathbf{V}_k^T; \\ \mathbf{T} &:= \mathbf{H} \mathbf{P}_{\mathcal{V}_k} \mathbf{L}^+, \end{aligned} \tag{A.3}$$

where  $\mathbf{P}_S$  denotes the orthogonal projection matrix onto  $S$ , and  $^+$  denotes the Moore–Penrose pseudo-inverse. Then for  $k \leq \text{rank}(\mathbf{L})$ ,

$$\mathbf{H} - \mathbf{T} \mathbf{L} = \mathbf{H} \mathbf{P}_{\mathcal{V}_k^\perp}. \tag{A.4}$$

**Proof.** Let  $\mathcal{N}(\mathbf{L})$  denote the nullspace of  $\mathbf{L}$ . Note that

$$\mathbf{L}^+ \mathbf{L} = \mathbf{P}_{\mathcal{N}(\mathbf{L})^\perp}.$$

Since  $k \leq \text{rank}(\mathbf{L})$ , then  $\mathcal{V}_k \subset \mathcal{N}(\mathbf{L})^\perp$ . Then from (A.3), it follows that

$$\begin{aligned} \mathbf{H} - \mathbf{T} \mathbf{L} &= \mathbf{H} - \mathbf{H} \mathbf{P}_{\mathcal{V}_k} \mathbf{L}^+ \mathbf{L}, \\ &= \mathbf{H} (\mathbf{I} - \mathbf{P}_{\mathcal{V}_k} \mathbf{P}_{\mathcal{N}(\mathbf{L})^\perp}) \\ &= \mathbf{H} (\mathbf{I} - \mathbf{P}_{\mathcal{V}_k}), \\ &= \mathbf{H} \mathbf{P}_{\mathcal{V}_k^\perp}. \quad \square \end{aligned}$$

This result allows us to bound the two key quantities,  $\|\mathbf{E}\| = \|\mathbf{H} - \mathbf{T} \mathbf{L}\|$  and  $\|\mathbf{T}\|$  in terms of  $\epsilon(\tau)$ .

**Lemma 4.** Let  $\tau$  and  $\epsilon(\tau)$  satisfy (A.1), and let  $\mathbf{T}$  be as defined in (A.3). Then, for  $\sigma_k$ , the  $k$ th largest singular value of  $\mathbf{L}$ ,

$$\|\mathbf{E}\|^2 = \|\mathbf{H} - \mathbf{T} \mathbf{L}\|^2 \leq \tau \sigma_{k+1}^2 + \epsilon(\tau), \tag{A.5}$$

$$\|\mathbf{T}\|^2 \leq \tau + \epsilon(\tau) \sigma_k^{-2}. \tag{A.6}$$

**Proof.** For any  $\mathbf{x} \in \mathbb{R}^N$  with  $\|\mathbf{x}\| \leq 1$ ,

$$\begin{aligned} \|(\mathbf{H} - \mathbf{T} \mathbf{L}) \mathbf{x}\|^2 &= \|\mathbf{H} \mathbf{P}_{\mathcal{V}_k^\perp} \mathbf{x}\|^2, \\ &\leq \tau \|\mathbf{L} \mathbf{P}_{\mathcal{V}_k^\perp} \mathbf{x}\|^2 + \epsilon(\tau) \|\mathbf{P}_{\mathcal{V}_k^\perp} \mathbf{x}\|^2, \\ &\leq \tau \sigma_{k+1}^2 + \epsilon(\tau), \end{aligned}$$

which proves (A.5). We have constructed  $\mathbf{T}$  so that

$$\mathbf{T} \mathbf{y} = \mathbf{0}, \quad \mathbf{y} \in \mathcal{R}(\mathbf{L})^\perp; \tag{A.7}$$

$$\mathbf{T} \mathbf{L} \mathbf{y} = \mathbf{0}, \quad \mathbf{y} \in \mathcal{V}_k^\perp, \tag{A.8}$$

where  $\mathcal{R}(\mathbf{L})$  is the range of  $\mathbf{L}$ . It follows that

$$\begin{aligned} \|\mathbf{T}\|^2 &= \max_{\|\mathbf{x}\| \leq 1} \|\mathbf{T} \mathbf{x}\|^2, \\ &\stackrel{(A.7)}{=} \max_{\|\mathbf{L} \mathbf{y}\| \leq 1} \|\mathbf{T} \mathbf{L} \mathbf{y}\|^2, \\ &\stackrel{(A.8)}{=} \max_{\|\mathbf{L} \mathbf{y}\| \leq 1, \mathbf{y} \in \mathcal{V}_k} \|\mathbf{H} \mathbf{y}\|^2, \\ &\leq \tau \|\mathbf{L} \mathbf{y}\|^2 + \epsilon(\tau) \|\mathbf{y}\|^2, \end{aligned}$$

for some  $\mathbf{y}$  achieving the maximum, and the last inequality uses the definition of  $\epsilon(\tau)$ . Note that  $\|\mathbf{L} \mathbf{y}\| \leq 1$ , and for  $\mathbf{y} \in \mathcal{V}_k$ ,  $\|\mathbf{y}\| \leq \sigma_k^{-1} \|\mathbf{L} \mathbf{y}\| \leq \sigma_k^{-1}$ . Therefore,

$$\|\mathbf{T}\|^2 \leq \tau + \epsilon(\tau) \sigma_k^{-2},$$

which proves (A.6).  $\square$

The proof of Theorem 1 follows from the combination of the bound provided in (8) and the results of Lemma 4. To explain the minimization of (11) over  $k$ , we note that the matrix  $\mathbf{T}$  in (3) can be constructed for any rank  $k$  not larger than the rank of the low-fidelity data matrix  $\mathbf{L}$ .

As a brief remark we note that (11) may instead be taken as

$$\|\mathbf{H} - \hat{\mathbf{H}}\| \leq \min_{\tau_1, \tau_2, k \leq \text{rank}(\mathbf{L})} (1 + \|\mathbf{C}_L\|) \sqrt{\tau_1 \sigma_{k+1}^2 + \epsilon(\tau_1)} + \|\mathbf{L} - \hat{\mathbf{L}}\| \sqrt{\tau_2 + \epsilon(\tau_2) \sigma_k^{-2}}, \quad (\text{A.9})$$

where the choices  $(\tau_1, \epsilon_1(\tau_1))$  and  $(\tau_2, \epsilon_2(\tau_2))$  both satisfy (A.1). Choosing different values may produce more effective values depending on  $\sigma_k$ ,  $\sigma_{k+1}$ ,  $\|\mathbf{C}_L\|$ , and  $\|\mathbf{L} - \hat{\mathbf{L}}\|$ . For simplicity, we have restricted the results of this work to the optimization over a single point  $(\tau, \epsilon(\tau))$ .

## References

- [1] M. Kennedy, A. O'Hagan, Predicting the output from a complex computer code when fast approximations are available, *Biometrika* 87 (1) (2000) 1–13.
- [2] Z. Qian, C. Seepersad, V. Joseph, J.K. Allen, C. Wu, Building surrogate models based on detailed and approximate simulations, *J. Mech. Des.* 128 (4) (2006) 668–677.
- [3] A. Forrester, A. Söbester, A. Keane, Multi-fidelity optimization via surrogate modelling, *Proc. R. Soc., Math. Phys. Eng. Sci.* 2088 (2007) 3251–3269.
- [4] J. Laurenceau, P. Sagaut, Building efficient response surfaces of aerodynamic functions with kriging and cokriging, *AIAA J.* 46 (2) (2008) 498–507.
- [5] W. Kleiber, S.R. Sain, M.J. Heaton, M. Wiltberger, C.S. Reese, D. Bingham, et al., Parameter tuning for a multi-fidelity dynamical model of the magnetosphere, *Ann. Appl. Stat.* 7 (2013) 1286–1310.
- [6] L. Le Gratiet, J. Garnier, Recursive co-kriging model for design of computer experiments with multiple levels of fidelity, *Int. J. Uncertain. Quantificat.* 4 (2014) 365–386.
- [7] L. Le Gratiet, C. Cannamela, Cokriging-based sequential design strategies using fast cross-validation techniques for multi-fidelity computer codes, *Technometrics* 57 (2015) 418–427.
- [8] P. Perdikaris, D. Venturi, J.O. Roysset, G.E. Karniadakis, Multi-fidelity modelling via recursive co-kriging and Gaussian–Markov random fields, *Proc. R. Soc., Math. Phys. Eng. Sci.* 471 (2015).
- [9] P. Perdikaris, D. Venturi, G.E. Karniadakis, Multifidelity information fusion algorithms for high-dimensional systems and massive data sets, *SIAM J. Sci. Comput.* 38 (2016) B521–B538.
- [10] L. Parussini, D. Venturi, P. Perdikaris, G.E. Karniadakis, Multi-fidelity Gaussian process regression for prediction of random fields, *J. Comput. Phys.* 336 (2017) 36–50.
- [11] M. Eldred, Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design, *AIAA Pap.* 2274 (2009) (2009) 37.
- [12] L.-T. Ng, M. Eldred, Multifidelity uncertainty quantification using nonintrusive polynomial chaos and stochastic collocation, in: *Proceedings of the 14th AIAA Non-Deterministic Approaches Conference*, Honolulu, HI, vol. 45, 2012, Number AIAA-2012-1852.
- [13] P. Palar, T. Tsuchiya, G. Parks, Decomposition-based evolutionary aerodynamic robust optimization with multi-fidelity point collocation non-intrusive polynomial chaos, in: *17th AIAA Non-Deterministic Approaches Conference*, 2015, p. 1377.
- [14] A.S. Padron, J.J. Alonso, M.S. Eldred, Multi-fidelity methods in aerodynamic robust optimization, in: *18th AIAA Non-Deterministic Approaches Conference*, 2016, p. 0680.
- [15] S. Heinrich, Multilevel Monte Carlo methods, in: *Large-Scale Scientific Computing*, Springer, 2001, pp. 58–67.
- [16] M. Giles, Multilevel Monte Carlo path simulation, *Oper. Res.* 56 (3) (2008) 607–617.
- [17] K. Cliffe, M. Giles, R. Scheichl, A. Teckentrup, Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients, *Comput. Vis. Sci.* 14 (1) (2011) 3–15.
- [18] A. Doostan, R. Ghanem, J. Red-Horse, Stochastic model reduction for chaos representations, *Comput. Methods Appl. Mech. Eng.* 196 (37–40) (2007) 3951–3966.
- [19] R. Ghanem, G. Saad, A. Doostan, Efficient solution of stochastic systems: application to the embankment dam problem, *Struct. Saf.* 29 (3) (2007) 238–251.
- [20] D. Kumar, M. Raisee, C. Lacor, An efficient non-intrusive reduced basis model for high dimensional stochastic problems in cfd, *Comput. Fluids* 138 (2016) 67–82.
- [21] A. Narayan, C. Gittelsohn, D. Xiu, A stochastic collocation algorithm with multifidelity models, *SIAM J. Sci. Comput.* 36 (2) (2014) A495–A521.
- [22] X. Zhu, A. Narayan, D. Xiu, Computational aspects of stochastic collocation with multifidelity models, *SIAM/ASA J. Uncertain. Quantificat.* 2 (1) (2014) 444–463.
- [23] B. Almroth, P. Stern, F. Brogan, Automatic choice of global shape functions in structural analysis, *AIAA J.* 16 (5) (1978) 525–528.
- [24] A. Noor, J. Peters, Reduced basis technique for nonlinear analysis of structures, *AIAA J.* 18 (4) (1980) 455–462.
- [25] J.S. Peterson, The reduced basis method for incompressible viscous flow calculations, *SIAM J. Sci. Stat. Comput.* 10 (1989) 777–786.
- [26] H. Cheng, Z. Gimbutas, P. Martinsson, V. Rokhlin, Reliable real-time solution of parametrized partial differential equations: reduced-basis output bound methods, *SIAM J. Sci. Comput.* 26 (4) (2005) 1389–1404.
- [27] G. Rozza, D. Huynh, T. Patera, Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations, *Arch. Comput. Methods Eng.* 15 (3) (2008) 229–275.
- [28] A. Doostan, G. Geraci, G. Iaccarino, A bi-fidelity approach for uncertainty quantification of heat transfer in a rectangular ribbed channel, in: *ASME Turbo Expo 2016: Turbomachinery Technical Conference and Exposition*, American Society of Mechanical Engineers, 2016, pp. V02CT45A031–V02CT45A031.
- [29] R.W. Skinner, A. Doostan, E.L. Peters, J.A. Evans, K.E. Jansen, An evaluation of bi-fidelity modeling efficiency on a general family of NACA airfoils, in: *35th AIAA Applied Aerodynamics Conference*, 2017, p. 3260.
- [30] H.R. Fairbanks, A. Doostan, C. Ketelsen, G. Iaccarino, A low-rank control variate for multilevel Monte Carlo simulation of high-dimensional uncertain systems, *J. Comput. Phys.* 341 (2017) 121–139.
- [31] L. Jofre, G. Geraci, H.R. Fairbanks, A. Doostan, G. Iaccarino, Multi-Fidelity Uncertainty Quantification of Irradiated Particle-Laden Turbulence, *Tech. Rep., Annual Research Brief, Center for Turbulence Research, Stanford University*, 2017.
- [32] A. Pinkus, *N-Widths in Approximation Theory*, Springer, 1985.
- [33] M. Bachmayr, A. Cohen, Kolmogorov widths and low-rank approximations of parametric elliptic PDEs, *Math. Comput.* 86 (304) (2017) 701–724, <https://doi.org/10.1090/mcom/3132>.
- [34] H. Cheng, Z. Gimbutas, P. Martinsson, V. Rokhlin, On the compression of low rank matrices, *SIAM J. Sci. Comput.* 26 (4) (2005) 1389–1404.
- [35] P. Martinsson, V. Rokhlin, M. Tytgert, A randomized algorithm for the decomposition of matrices, *Appl. Comput. Harmon. Anal.* 30 (1) (2011) 47–68.

- [36] N. Halko, P.-G. Martinsson, J.A. Tropp, Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Rev.* 53 (2) (2011) 217–288.
- [37] G.H. Golub, C.F. Van Loan, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, MD, USA, 1996.
- [38] M. Mahoney, P. Drineas, Cur matrix decompositions for improved data analysis, *Proc. Natl. Acad. Sci.* 106 (2009) 697–702.
- [39] E. Elhamifar, R. Vidal, Sparse subspace clustering, in: *IEEE Conference on Computer Vision and Pattern Recognition, 2009, CVPR 2009, IEEE, 2009*, pp. 2790–2797.
- [40] E. Dyer, T. Goldstein, R. Patel, K. Kording, R. Baraniuk, Self-expressive decompositions for matrix approximation and clustering, arXiv preprint arXiv:1505.00824.
- [41] D.J. Perry, R.M. Kirby, A. Narayan, R.T. Whitaker, Allocation strategies for high fidelity models in the multifidelity regime, submitted for publication.
- [42] D. Xiu, J. Hesthaven, High-order collocation methods for differential equations with random inputs, *SIAM J. Sci. Comput.* 27 (3) (2005) 1118–1139.
- [43] S. Hosder, R. Walters, R. Perez, A non-intrusive polynomial chaos method for uncertainty propagation in CFD simulations, in: *44th AIAA Aerospace Sciences Meeting and Exhibit, AIAA-2006-891, Reno (NV), 2006*.
- [44] A. Doostan, H. Owhadi, A non-adapted sparse approximation of PDEs with stochastic inputs, *J. Comput. Phys.* 230 (2011) 3015–3034.
- [45] P.G. Constantine, M. Eldred, E. Phipps, Sparse pseudospectral approximation method, *Comput. Methods Appl. Mech. Eng.* 229 (2012) 1–12.
- [46] M. Hadigol, A. Doostan, Least squares polynomial chaos expansion: a review of sampling strategies, *Comput. Methods Appl. Mech. Eng.* 332 (2018) 382–407.
- [47] T. Bui-Thanh, K. Willcox, O. Ghattas, Model reduction for large-scale systems with high-dimensional parametric input space, *SIAM J. Sci. Comput.* 30 (6) (2008) 3270–3288.
- [48] J. Jiang, Y. Chen, A. Narayan, Offline-enhanced reduced basis method through adaptive construction of the surrogate training set, *J. Sci. Comput.* 73 (2–3) (2017) 853–875.
- [49] O. LeMaitre, M. Reagan, H. Najm, R. Ghanem, O. Knio, A stochastic projection method for fluid flow. II: Random process, *J. Comput. Phys.* 181 (2002) 9–44.
- [50] O.L. Maitre, O. Knio, *Spectral Methods for Uncertainty Quantification with Applications to Computational Fluid Dynamics*, Springer, 2010.
- [51] P.L. Quéré, Accurate solutions to the square thermally driven cavity at high Rayleigh number, *Comput. Fluids* 20 (1) (1991) 29–41.
- [52] J. Peng, J. Hampton, A. Doostan, A weighted  $\ell_1$ -minimization approach for sparse polynomial chaos expansions, *J. Comput. Phys.* 267 (2014) 92–111.
- [53] J. Hampton, A. Doostan, Compressive sampling of polynomial chaos expansions: convergence analysis and sampling strategies, *J. Comput. Phys.* 280 (2015) 363–386.
- [54] J. Hampton, A. Doostan, Coherence motivated sampling and convergence analysis of least squares polynomial chaos regression, *Comput. Methods Appl. Mech. Eng.* 290 (2015) 73–97.
- [55] J. Hampton, A. Doostan, Basis adaptive sample efficient polynomial chaos (base-pc), arXiv preprint arXiv:1702.01185.
- [56] A. Logg, K.-A. Mardal, G.N. Wells, et al., *Automated Solution of Differential Equations by the Finite Element Method*, Springer, 2012.