

ARTICLE TYPE

A Generalized Probabilistic Learning Approach for Multi-Fidelity Uncertainty Propagation in Complex Physical Simulations

J. Nitzler^{1,2} | J. Biehler¹ | N. Fehn¹ | P.-S. Koutsourelakis² | W. A. Wall¹¹Institute for Computational Mechanics,
Technical University of Munich, Germany²Professorship of Continuum Mechanics,
Technical University of Munich, Germany**Correspondence**

*W. A. Wall, Email: wall@lnm.mw.tum.de

Present AddressBoltzmannstrasse 15, 85748 Garching b.
München**Abstract**

Two of the most significant challenges in uncertainty propagation pertain to the high computational cost for the simulation of complex physical models and the high dimension of the random inputs. In applications of practical interest both of these problems are encountered and standard methods for uncertainty quantification either fail or are not feasible. To overcome the current limitations, we propose a probabilistic multi-fidelity framework that can exploit lower-fidelity model versions of the original problem in a *small data* regime. The approach circumvents the curse of dimensionality by learning dependencies between the outputs of high-fidelity models and lower-fidelity models instead of explicitly accounting for the high-dimensional inputs. We complement the information provided by a low-fidelity model with a low-dimensional set of informative features of the stochastic input, which are discovered by employing a combination of supervised and unsupervised dimensionality reduction techniques. The goal of our analysis is an efficient and accurate estimation of the full probabilistic response for a high-fidelity model. Despite the incomplete and noisy information that low-fidelity predictors provide, we demonstrate that accurate and certifiable estimates for the quantities of interest can be obtained in the *small data* regime, i.e., with significantly fewer high-fidelity model runs than state-of-the-art methods for uncertainty propagation. We illustrate our approach by applying it to challenging numerical examples such as Navier-Stokes flow simulations and monolithic fluid-structure interaction problems.

KEYWORDS:

Uncertainty Quantification; Probabilistic Learning; Multi-Fidelity Schemes; Bayesian Inference; Fluid-Structure Interaction; Discontinuous Galerkin Method; Inverse Problems; High-Dimensions

1 | INTRODUCTION

The analysis of complex real-world systems is usually based on sophisticated, high-fidelity (HF) computer models. Accuracy comes at the cost of computationally expensive simulations, characterized by detailed physical resolution, fine temporal and spatial discretizations, as well as narrow numerical tolerances. A single evaluation of such models, for example large scale nonlinear and transient biomechanical problems or coupled fluid simulations, can take hours or days even on modern high-performance clusters. Nevertheless, many questions in industry and science require a vast number of accurate computer simulations to understand different system configurations, boundary conditions, perform optimization tasks, or investigate forward and backward

uncertainty propagation. Unfortunately, limitations in available resources render the aforementioned types of analysis unfeasible, so that for most practical applications, analysts either avoid such fundamental investigations completely, or fall back to less accurate but cheaper low-fidelity (LF) variants of the original problem to conduct the analysis.

One strategy to overcome these problems pertains to multi-fidelity schemes which, by combining information provided by different levels of model sophistication, attempt to decrease the number of high-fidelity model runs required, while retaining the same accuracy^{1,2}. Especially sampling based methods for uncertainty propagation, while often being the only choice for nonlinear problems with large variabilities, become unfeasible for costly numerical models. Multi-level Monte Carlo methods (MLMC)^{3,4,5,6} were some of the earliest schemes used to accelerate the calculation of the expectation and variance of a quantity of interest (QoI) on complex models, given uncertain inputs \mathbf{x} . Unfortunately, the method necessitates linear dependence between model outputs and can only yield asymptotic error estimates. The estimation of the whole response distribution is restricted to special cases^{7,6}. Other contributions used low-fidelity model versions to identify important regions in the input space, motivating adaptive sampling strategies and multi-fidelity importance sampling schemes^{8,9,10}, while still requiring costly sampling of the HF model. Similar ideas arose for inverse problems in the form of multi-stage Markov-chain Monte Carlo methods^{11,12,13}. Promising alternative approaches that recently gained considerable attention are the so-called Bayesian multi-fidelity schemes^{14,15,16,17,18,19,20,21,22,23,24,25}. We already demonstrated their superior performance for large scale numerical problems²².

Thus far, most state-of-the-art multi-fidelity methods aim to construct an approximation for the high-fidelity (HF) computer simulation $y_{\text{HF}}(\mathbf{x})$ in the form of a surrogate model $\hat{f}(\mathbf{x}) \approx y_{\text{HF}}(\mathbf{x})$ that is built based on a *small* number of HF simulations. Such multi-fidelity approaches exploit the characteristic that the discrepancy of HF and LF model responses has a simpler mathematical structure than the HF model response $y_{\text{HF}}(\mathbf{x})$ itself and hence, the error over \mathbf{x} can be efficiently learned to yield good HF response predictions using few data^{16,17,23,24,26}. Unfortunately, such approaches face serious problems for applications with high stochastic dimensionality, especially in the case of a *small data* scenario which refers to the small amount of HF simulations available due to the associated high costs.

This contribution expands upon our previous work^{14,22} in presenting a generalized multi-fidelity framework that is well suited for uncertainty propagation in very high stochastic dimensions. The advances proposed lead to a higher accuracy so that we can exploit a wider range of automatically generated low-fidelity versions of the original problem. We provide a general theoretic viewpoint on the topic of Bayesian uncertainty propagation while emphasizing the practical applicability of the proposed techniques towards a broad field of engineering problems. Bayesian multi-fidelity uncertainty propagation is a non-intrusive, data-driven approach that can be used with any numerical solver. In general, only one low-fidelity model is required but the method can seamlessly integrate several sources of information to improve its performance. The principal result of the method is a full probabilistic description of the quantity of interest (QoI) in the form of an approximation of the high-fidelity probability density $p(y_{\text{HF}})$. Point estimates such as event probabilities, expectations or maximum likelihood estimates can then be calculated along with credible intervals due to the Bayesian nature of our approach.

The paper is structured as followed: In Section 2, we present the theoretical foundation for Bayesian forward and backward uncertainty propagation, as well as the key idea for the multi-fidelity extension of both cases. We then discuss the challenges associated with a small data scenario and introduce a probabilistic learning strategy that leads to accurate prediction of the HF model's output density by exploiting the information provided by the output of a LF model as well as by informative features of the model's input. We demonstrate how such informative features of the model input can be learned based on a combination of supervised and unsupervised dimensionality reduction techniques, at negligible extra cost. The derivations are first formulated for any probabilistic regression model before they are specialized for Gaussian Processes. We demonstrate that our Bayesian multi-fidelity approach can provide credible intervals for the density estimate itself. As we aim at complex physical simulations, we additionally present strategies for the automated generation of efficient LF model versions that ultimately lead to significant computational speed-ups of our multi-fidelity approach, compared to other state-of-the-art methods for uncertainty quantification. The steps of the generalized Bayesian multi-fidelity approach for uncertainty quantification are summarized in a pseudo-algorithm in Section 2.4. Section 3 provides theoretical suggestions for an automated creation of low-fidelity model versions based on numerical relaxation. Afterwards, the framework is discussed and demonstrated for challenging fluid and fluid-structure interaction problems in Section 4. Apart from the algorithmic aspects, we also focus on the modeling of physically compliant random boundary conditions, random fields, and their numerical realization. We conclude with a discussion of the numerical results and computational performance, and provide an outlook on possible future developments.

2 | BAYESIAN FORWARD AND BACKWARD UNCERTAINTY PROPAGATION

Uncertainty quantification (UQ) can be distinguished into forward and backward uncertainty propagation (the latter being known as inverse problems or model calibration problems). Forward UQ aims to propagate the uncertainty of a random input vector \mathbf{x} with a given density $p(\mathbf{x})$ through a physics-based, high-fidelity, numerical model to accurately and efficiently quantify the uncertainty of one or more outputs or Quantities of Interest (QoIs) \mathbf{y} , for example in the form of their density $p(\mathbf{y}_{\text{HF}})$. The random vector \mathbf{x} can represent uncertainties in model parameters, loads, excitations or boundary and initial conditions. For applications of practical interest its dimension is very high (in the hundreds or thousands). On the other hand, in backward uncertainty propagation, given a similar mathematical model and, in general, noisy observations Y_{obs} of the system's output \mathbf{y} , the goal is to estimate a vector of model inputs \mathbf{x} .

In the following we denote by $y_{\text{HF}}(\mathbf{x})$ the *deterministic* input-output map implied by a high-fidelity model which in most cases of practical interest is not available in closed form and expensive to evaluate (e.g., for each value of \mathbf{x} the numerical solution of time-dependent, nonlinear PDEs needs to be carried out). We assume that the high-fidelity model is the reference model, i.e., its predictions y_{HF} coincide with the QoI. For clarity of the presentation we consider the scalar case, i.e. $y_{\text{HF}} : \mathbb{R}^d \rightarrow \mathbb{R}$, where $d = \dim(\mathbf{x})$. Furthermore, to simplify the notation, we make no distinction between random variables and the values these can take. In this notation, \mathbf{x} or y_{HF} denote the respective random variables and possible realizations, whereas $y_{\text{HF}}(\mathbf{x})$ refers to a deterministic function. Scalar quantities are expressed by plain letters (e.g., y_{HF} for a scalar, high-fidelity model output), in contrast to boldface letters (such as the input vector \mathbf{x}), which denote vector-valued quantities. We denote with capital letters a data set that can either consist of scalar quantities or vector-valued quantities. Data sets of scalar quantities are written in plain capital letters such as the vector of row-wise scalar experimental observations Y_{obs} . On the other hand, vector-valued quantities, such as the matrix of row-wise vector-valued model inputs \mathbf{X} , are written with boldface capital letters. The most important distinction is made between the training data (indicated by capital letter but without further superscripts, e.g., Y_{HF}) and test data that has an asterisk superscript (e.g., \mathbf{x}^* for one arbitrary test input or the large data set of all test inputs \mathbf{X}^*).

For forward UQ we seek the whole response density $p(y_{\text{HF}})$ which can then be used to calculate any statistic of interest. The resulting output density for the QoI can be expressed as the integral over the conditional distribution $p(y_{\text{HF}}|\mathbf{x})$ weighted by the density of the input $p(\mathbf{x})$. In the special case of a deterministic function $y_{\text{HF}}(\mathbf{x})$, the conditional distribution $p(y_{\text{HF}}|\mathbf{x})$ can be expressed in form of a Dirac distribution $p(y_{\text{HF}}|\mathbf{x}) = \delta_{y_{\text{HF}}}(y_{\text{HF}} - y_{\text{HF}}(\mathbf{x}))$:

$$p(y_{\text{HF}}) = \int_{\Omega_{\mathbf{x}}} p(y_{\text{HF}}|\mathbf{x}) p(\mathbf{x}) d\mathbf{x} = \int_{\Omega_{\mathbf{x}}} \delta_{y_{\text{HF}}}(y_{\text{HF}} - y_{\text{HF}}(\mathbf{x})) p(\mathbf{x}) d\mathbf{x} \quad (1)$$

Equation (1) is usually approximated by Monte Carlo methods which, depending on the demanded level of accuracy and especially for the tails of $p(y_{\text{HF}})$ (i.e., rare events), would require a huge amount of evaluations of $y_{\text{HF}}(\mathbf{x})$. The overall computational cost can render such an approach impracticable or unfeasible. Alternative strategies have attempted to approximate the map $y_{\text{HF}}(\mathbf{x})$ or the conditional density $p(y_{\text{HF}}|\mathbf{x})$ using a variety of surrogates or emulators which are generally trained on n_{train} simulation data pairs, i.e., $\mathcal{D}_{\text{HF}} = \{\mathbf{x}_i, y_{\text{HF}}(\mathbf{x}_i)\}_{i=1}^{n_{\text{train}}}$. Given the high dimension d of \mathbf{x} , this task gives rise to several accuracy and efficiency challenges. Even when the most expressive, modern machine learning tools are deployed (e.g., Deep Neural Nets) the number n_{train} of high-fidelity evaluations needed to achieve an acceptable level of accuracy can render such methods impracticable or unfeasible as well.

For inverse problems, the answer in the Bayesian framework is provided in the form of the so-called posterior probability density on \mathbf{x} , given the observations, i.e. $p(\mathbf{x}|Y_{\text{obs}})$, which arises as a combination of the prior density $p(\mathbf{x})$ and the likelihood:

$$\underbrace{p(\mathbf{x}|Y_{\text{obs}})}_{\text{Posterior}} \propto \underbrace{p(Y_{\text{obs}}|y_{\text{HF}}(\mathbf{x}))}_{\text{Likelihood}} \underbrace{p(\mathbf{x})}_{\text{Prior}} \quad (2)$$

From the posterior we can compute point estimates \mathbf{x}_{est} that represent the vector of calibrated model parameters. The Bayesian interpretation of backward uncertainty propagation provides a natural regularization framework for ill-posed problems. The computationally expensive part is the likelihood $p(Y_{\text{obs}}|y_{\text{HF}}(\mathbf{x}))$, as for each evaluation, the HF model must be solved. Usually, the posterior distribution $p(\mathbf{x}|Y_{\text{obs}})$ is approximated by advanced sampling methods such as Markov-Chain Monte Carlo (MCMC) sampling or sequential Monte-Carlo (SMC) methods. Nevertheless, even modern sampling techniques also require an unfeasible amount of costly HF model evaluations.

In this contribution, we will focus on an efficient approach for forward uncertainty propagation but also highlight the analogies to Bayesian model calibration. An in-depth discussion of the latter case is, however, outside the scope of this paper.

2.1 | General Aspects of Bayesian Multi-Fidelity Forward and Inverse Problems

The previous expressions (1) and (2) involved a computationally expensive high-fidelity computer model implied by $y_{\text{HF}}(\mathbf{x})$. In the following, we demonstrate how less expensive lower-fidelity models in combination with low-dimensional features of \mathbf{x} can be employed to obtain accurate and certifiable estimates of the aforementioned quantities requiring only very small numbers of high-fidelity runs. In the simplest version, we presuppose the availability of a lower-fidelity model which provides a potentially very poor approximation of the QoI. In the scalar case, we denote this with y_{LF} and the associated input-output (deterministic) map by $y_{\text{LF}}(\mathbf{x})$.

In contrast to multi-level Monte-Carlo techniques, which also make use of lower-fidelity models in combination with frequentist estimators, we advocate a Bayesian perspective¹⁴, which we refer to as *Bayesian Multi-Fidelity Monte-Carlo* (BMFMC) method^{22,18}. The basis of the framework advocated is re-expressing the sought density as:

$$\begin{aligned}
 p(y_{\text{HF}}) &= \int_{\Omega_{\mathbf{x}}} \underbrace{p(y_{\text{HF}}|\mathbf{x})}_{\substack{\text{Dirac:} \\ \text{comp. expensive}}} \cdot p(\mathbf{x}) \, d\mathbf{x} && \rightarrow \text{standard forward UQ, see equation (1)} \\
 &= \int_{\Omega_{y_{\text{LF}}}} \int_{\Omega_{\mathbf{x}}} p(y_{\text{HF}}, y_{\text{LF}}, \mathbf{x}) \, d\mathbf{x} dy_{\text{LF}} && \rightarrow \text{write as joint density and expand expression by LF model response } y_{\text{LF}} \\
 &= \int_{\Omega_{y_{\text{LF}}}} \int_{\Omega_{\mathbf{x}}} p(y_{\text{HF}}, \mathbf{x}|y_{\text{LF}}) \cdot p(y_{\text{LF}}) \, d\mathbf{x} dy_{\text{LF}} && \rightarrow \text{condition on LF model } y_{\text{LF}} \\
 &= \int_{\Omega_{y_{\text{LF}}}} \underbrace{p(y_{\text{HF}}|y_{\text{LF}})}_{\substack{\text{Approximate with} \\ \text{little HF data}}} \cdot \underbrace{p(y_{\text{LF}})}_{\substack{\text{Sampling} \\ \text{on LF}}} \, dy_{\text{LF}} && \rightarrow \text{integrate over inputs } \mathbf{x} \text{ to yield multi-fidelity formulation in BMFMC}^{22}
 \end{aligned} \tag{3}$$

We note that none of the expressions above contain *any errors or approximations*. Furthermore, the crucial conditional density $p(y_{\text{HF}}|y_{\text{LF}})$ that must be learned or estimated is *independent of the dimension of the input vector \mathbf{x}* . The premise of *BMFMC* is that all the densities above can be estimated at a cost (as measured by the number of high-fidelity solves) which is much less than the alternatives. In the case of $p(y_{\text{LF}})$ this can be achieved as long as the lower-fidelity model is much cheaper than the high-fidelity reference. To assess the feasibility of this task for $p(y_{\text{HF}}|y_{\text{LF}})$ and to better understand the role of this conditional density we consider the following limiting cases:

extreme 1) The lower-fidelity model y_{LF} is independent of y_{HF} , i.e., $p(y_{\text{HF}}|y_{\text{LF}}) = p(y_{\text{HF}})$. While (3) remains valid, any attempt to estimate $p(y_{\text{HF}}|y_{\text{LF}})$ will be comparable to a Monte Carlo estimator applied directly on y_{HF} . Hence, it is unlikely that any significant efficiency gains could be achieved.

extreme 2) The lower- and high-fidelity model are fully dependent, i.e., there is a function, say f , such that $y_{\text{HF}} = f(y_{\text{LF}})$ and $p(y_{\text{HF}}|y_{\text{LF}}) = \delta_{y_{\text{HF}}}(y_{\text{HF}} - f(y_{\text{LF}}))$. Any efficiency gains, in this case, would depend on the probabilistic machine learning model postulated for learning the conditional.

Given a physically motivated $y_{\text{LF}}(\mathbf{x})$, one would expect the actual $p(y_{\text{HF}}|y_{\text{LF}})$ to be between these two extremes as shown in Figure 1. Errors will only be introduced through the numerical approximation of the densities in equation (3). We distinguish therefore between the following sources of error:

error 1) The primary source of error in the method stems from the probabilistic model (e.g., Gaussian Processes or other probabilistic regression tools) selected to approximate the conditional distribution $p(y_{\text{HF}}|y_{\text{LF}})$. If the family of approximating densities considered does not include the true one, a modeling error will be introduced. For a Gaussian Process model in $\Omega_{y_{\text{HF}} \times y_{\text{LF}}}$, such assumptions are, i.e., a normal distributed conditional $p(y_{\text{HF}}|y_{\text{LF}})$ along with smoothness properties of the random process, imposed by the selected kernel.

error 2) The second error source pertains to the amount of available training data (i.e., pairs of lower- and high-fidelity runs). Even if the true $p(y_{\text{HF}}|y_{\text{LF}})$ belongs to the selected model class, it is unlikely that it would be recovered exactly with limited data.

BMFMC brings two major advantages for UQ with computationally demanding computer models: Firstly, by exploiting information encoded in computationally cheaper, lower-fidelity (LF) versions of the original computer model, it can drastically reduce the number of costly HF model evaluations and enable forward and backward uncertainty propagation even for very expensive models. Secondly, by learning the statistical dependence between the y_{LF} and y_{HF} , *BMFMC* circumvents the curse of dimensionality which arises as a result of high-dimensional model inputs \mathbf{x} . As an explicit treatment of the high-dimensional input vector \mathbf{x} in surrogates is not expedient, especially for a low number of HF data, we make use instead of statistical dependencies between y_{LF} and y_{HF} .

Figure 1 provides a visualization of the main terms in equation (3). In particular, Figure 1(a) shows examples of response surfaces for a two-dimensional input \mathbf{x} . The upper response surface represents a LF model and the lower one the corresponding HF model response. A red dot marks a function value for the same \mathbf{x} on both models. The corresponding Dirac density $p(y_{HF}|\mathbf{x})$

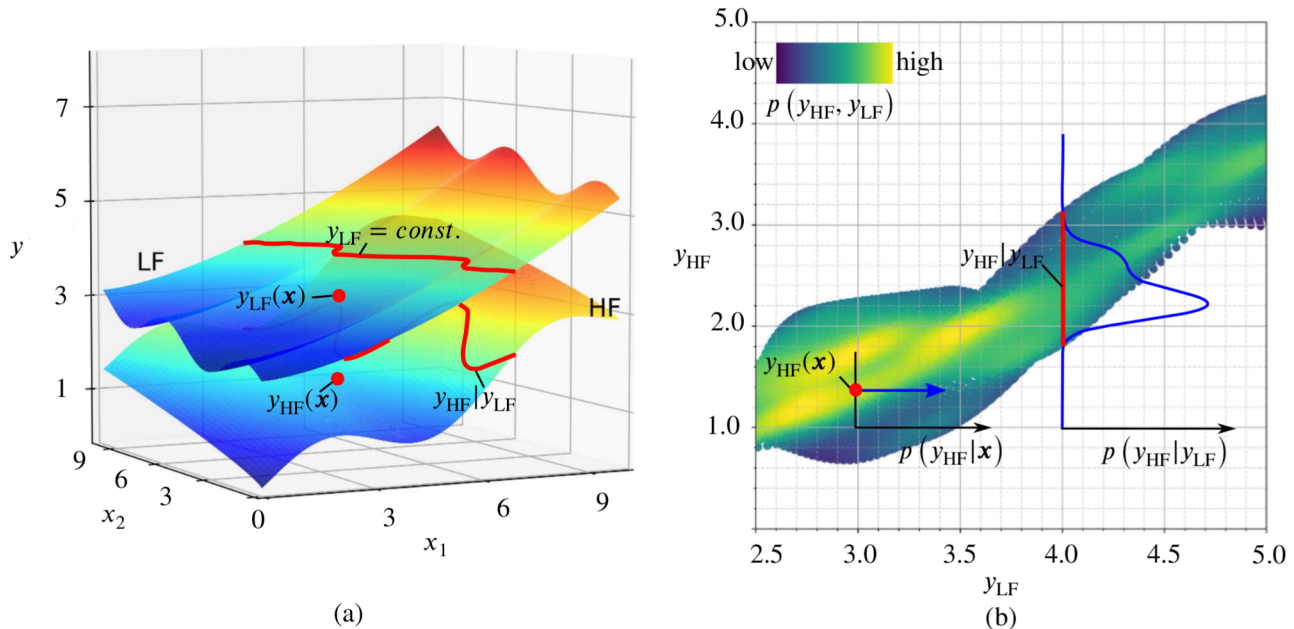


FIGURE 1 Visualization of HF and LF model dependencies. Left: Example of LF and HF model outputs with two input variables x_1 and x_2 . Right: Dependence between LF and HF output. The joint density $p(y_{LF}, y_{HF})$ is color-coded. Conditional densities $p(y_{HF}|\mathbf{x})$ and $p(y_{HF}|y_{LF})$ are shown as slices of $p(y_{LF}, y_{HF})$ in blue.

is shown by a blue arrow, centered on the red dot in Figure 1(b). An indicative conditional density of y_{HF} given y_{LF} is also shown in Figure 1(b). The vertical red line shows the support for the corresponding conditional density $p(y_{HF}|y_{LF})$ which encodes the knowledge about possible outcomes of y_{HF} when only y_{LF} is observed (without information of a specific \mathbf{x} that yielded y_{LF}).

In this contribution we generalize expression (3) by considering, next to LF models, also informative features $\boldsymbol{\gamma}(\mathbf{x}) = \boldsymbol{\gamma}_j(\mathbf{x})$ of the input \mathbf{x} , with $j \in \mathbb{N} : [1, N_\gamma]$ and N_γ being the number of LF models used in the multi-fidelity approach. We will further elaborate on this in Section 2.4. We denote the vector of informative features $\boldsymbol{\gamma}(\mathbf{x})$ and LF responses $y_{LF}(\mathbf{x})$ as $\mathbf{z}_{LF}(\mathbf{x}) = [y_{LF}(\mathbf{x}), \boldsymbol{\gamma}(\mathbf{x})]^T$. With \mathbf{z}_{LF} , we jointly denote the corresponding random vector as well as values that this can take. The basic elements of *BMFMC* remain unaltered if one employs multiple low-fidelity features, summarized in \mathbf{z}_{LF} , so that equation (3) becomes:

$$p(y_{HF}) = \int_{\Omega_{\mathbf{z}_{LF}}} p(y_{HF}|\mathbf{z}_{LF}) \cdot p(\mathbf{z}_{LF}) d\mathbf{z}_{LF} \quad (4)$$

We demonstrate in the subsequent sections how the modeling error (i.e., **error 1** above) can be reduced and superior estimates can be obtained by an appropriate selection of the input features.

In order to show the generality and power of the proposed strategy, we would like to conclude this section by giving a short outlook on the applicability of the Bayesian multi-fidelity approach for *inverse uncertainty propagation* which we derive from

the standard formulation in equation (2):

$$\begin{aligned}
p(\mathbf{x}|Y_{\text{obs}}) &\propto \underbrace{p(Y_{\text{obs}}|y_{\text{HF}}(\mathbf{x}))}_{\text{Likelihood}} \cdot p(\mathbf{x}) && \rightarrow \text{Standard Bayesian inverse problem} \\
&= \int_{\Omega_{y_{\text{HF}}}} p(Y_{\text{obs}}|y_{\text{HF}}) \cdot p(y_{\text{HF}}|\mathbf{x}) dy_{\text{HF}} \cdot p(\mathbf{x}) && \rightarrow \text{Expand the likelihood} \\
&\approx \int_{\Omega_{y_{\text{HF}}}} \underbrace{p(Y_{\text{obs}}|y_{\text{HF}}) \cdot p(y_{\text{HF}}|\mathbf{z}_{\text{LF}}(\mathbf{x}))}_{p(Y_{\text{obs}}|\mathbf{z}_{\text{LF}}(\mathbf{x}))} dy_{\text{HF}} \cdot p(\mathbf{x}) && \rightarrow \text{Approximation: } p(y_{\text{HF}}|\mathbf{z}_{\text{LF}}(\mathbf{x})) \approx p(y_{\text{HF}}|\mathbf{x}) \quad (5) \\
&= \int_{\Omega_{y_{\text{HF}}}} \int_{\Omega_{\mathbf{z}_{\text{LF}}}} p(Y_{\text{obs}}|y_{\text{HF}}) \cdot p(y_{\text{HF}}|\mathbf{z}_{\text{LF}}) && \rightarrow \text{Expand } p(y_{\text{HF}}|\mathbf{z}_{\text{LF}}(\mathbf{x})) \\
&\quad \cdot p(\mathbf{z}_{\text{LF}}|\mathbf{x}) d\mathbf{z}_{\text{LF}} dy_{\text{HF}} \cdot p(\mathbf{x}) \\
&:= p(\mathbf{x}|Y_{\text{obs}})_{\text{MF}}
\end{aligned}$$

In the equations above we approximate the computationally demanding $p(y_{\text{HF}}|\mathbf{x})$ with $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}}(\mathbf{x}))$. In contrast to the multi-fidelity forward UQ formulation, the multi-fidelity inverse problem contains an approximation step, by replacing $p(y_{\text{HF}}|\mathbf{x})$ with $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}}(\mathbf{x}))$. In contrast to \mathbf{x} , the vector $\mathbf{z}_{\text{LF}}(\mathbf{x})$ is low-dimensional so that $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}}(\mathbf{x}))$ can be well approximated with probabilistic learning methods. It is important to point out that in the multi-fidelity scenario, the HF output y_{HF} (for each \mathbf{x}) becomes an unobserved quantity which in fully-Bayesian fashion is modeled as a random variable. In the last line of equation (5) we expand the term $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}}(\mathbf{x}))$ to show the analogies to equation (4). The Bayesian multi-fidelity posterior $p(\mathbf{x}|Y_{\text{obs}})_{\text{MF}}$ reflects not only the uncertainties in the data but also the uncertainty introduced by \mathbf{z}_{LF} for predicting y_{HF} (instead of solving the HF model directly)¹⁵. In this contribution, the Bayesian multi-fidelity inverse problem should only illustrate the potential of our multi-fidelity framework. The approximation of the involved terms brings further technical challenges that are outside the scope of the current paper.

2.2 | Learning of $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}})$: Multi-Fidelity Forward UQ in the Small Data Case

In the preceding formulations, the densities involved were assumed to be known. In real applications, an analytic description of the density terms is not existent but only the *small data set* of HF model evaluations and the computationally cheaper *large data set* of LF features \mathbf{z}_{LF} is available. Classic density estimation techniques are inaccurate in small data scenarios. Hence, this section focuses on strategies to efficiently learn the multi-fidelity conditional distribution $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}})$ which is the key element of the multi-fidelity approach. Apart from the obvious accuracy requirements, it is essential that the necessary number of training data $\mathcal{D}_f = \{\mathbf{Z}_{\text{LF}}, Y_{\text{HF}}\}$ with $\mathbf{Z}_{\text{LF}} = \mathbf{z}_{\text{LF}}(\mathbf{X})$ and $Y_{\text{HF}} = y_{\text{HF}}(\mathbf{X})$, is minimized.

As one regression function $f : \mathbb{R}^{d_{\text{z}_{\text{LF}}}} \rightarrow \mathbb{R}$, with $d_{\text{z}_{\text{LF}}} = \dim(\mathbf{z}_{\text{LF}})$, for the data set \mathcal{D}_f would not describe the problem adequately, due to the noisy relationship between \mathbf{z}_{LF} and y_{HF} , we follow a probabilistic learning approach, considering a distribution of possible regression functions f that would describe the small data set \mathcal{D}_f and obtain a density estimate for the HF output. In the Bayesian setting the multi-fidelity distribution $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}})$ is approximated in two steps: First we define a prior distribution for possible regression functions $p(f|\mathbf{z}_{\text{LF}})$ denoted as the prior probabilistic model, which incorporates our beliefs (e.g., smoothness, dependency) about the relationship of y_{HF} and \mathbf{z}_{LF} . Since f is *infinite-dimensional*, we follow^{27,28} in the slight abuse of notation for denoting with $p(f)$ and $p(f|\mathcal{D}_f)$ the prior and posterior of f , respectively. Please note, that $p(f|\mathcal{D}_f)$ implies a random process, e.g., a Gaussian Process $\mathcal{GP}_f(m_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}), v_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}))$, while $p(f^*|\mathbf{z}_{\text{LF}}^*, \mathcal{D}_f)$ describes the *uni-variate* conditional distribution for a certain test value \mathbf{z}_{LF}^* which is a normal distribution $\mathcal{N}_{f^*}(m_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}^*), v_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}^*))$, in case a Gaussian Process was deployed as a probabilistic model. We write f^* to denote the uni-variate random variable that emerges for the evaluation of the random process at a particular test input \mathbf{z}_{LF}^* . We denote the value $p(y_{\text{HF}}^*|\mathbf{z}_{\text{LF}}^*, f^*)$, with $p(f^*|\mathbf{z}_{\text{LF}}^*, \mathcal{D}_f)$, in contrast to the former expression $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}})$:

$$p(y_{\text{HF}}|\mathbf{z}_{\text{LF}}) \rightarrow p(y_{\text{HF}}^*|\mathbf{z}_{\text{LF}}^*, f^*, \mathcal{D}_f) \text{ with } p(f^*|\mathbf{z}_{\text{LF}}^*, \mathcal{D}_f) \quad (6)$$

The regression function f is only needed to construct the probabilistic model and allows for a structured representation of our approach but we are not interested in particular values of the function itself²⁹. In the subsequent predictions for the HF model's output density, we will hence calculate statistics for the variability in f^* and eliminate the dependency on the latter.

The starting point of the multi-fidelity UQ is a sampling procedure on the LF model, resulting in a large LF data set $\mathcal{D}_{\text{LF}}^* = \{\mathbf{x}_i^*, y_{\text{LF}}(\mathbf{x}_i^*)\}_{i=1}^{N_{\text{sample}}} = \{\mathbf{X}^*, \mathbf{Y}_{\text{LF}}^*\}$. The sample size N_{sample} is dependent on the simulation model and the statistics of interest but usually at least in the hundreds or thousands. Based on $\mathcal{D}_{\text{LF}}^*$ and criteria that will be discussed in Section 2.4, we select a subset of optimal inputs $\mathbf{X} \subset \mathbf{X}^*$ with size $n_{\text{train}} \ll N_{\text{sample}}$ and run simulations that yield the corresponding HF outputs $\mathbf{Y}_{\text{HF}} = y_{\text{HF}}(\mathbf{X})$. Following the procedure described in Section 2.4, we can additionally learn informative features γ_i , *without the need of further simulation runs*. The tuples of corresponding HF and LF outputs and informative features yield the training data set $\mathcal{D}_f = \{\mathbf{Z}_{\text{LF}}, \mathbf{Y}_{\text{HF}}\} = \{\mathbf{z}_{\text{LF}}(\mathbf{X}), y_{\text{HF}}(\mathbf{X})\}$. The extended small data approximation of the multi-fidelity conditional distribution $p(y_{\text{HF}}^* | f^*, \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f)$ can now be plugged into equation (3) to yield the Bayesian multi-fidelity forward uncertainty estimate for the limited data case:

$$p(y_{\text{HF}}^* | f^*, \mathcal{D}_f) = \int_{\Omega_{\text{z}_{\text{LF}}}} \underbrace{p(y_{\text{HF}}^* | f^*, \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f)}_{\text{Likelihood of HF observations}} \cdot \underbrace{p(\mathbf{z}_{\text{LF}})}_{\text{Marginal density: direct MC on LF model}} d\mathbf{z}_{\text{LF}} \quad (7)$$

Now we can among others calculate the expectation and variance of equation (7) with respect to the random variable f^* . The expectation $\mathbb{E}_{f^*} [p(y_{\text{HF}}^* | f^*, \mathcal{D}_f)]$ serves as an approximation for the HF distribution $\mathbb{E}_{f^*} [p(y_{\text{HF}}^* | f^*, \mathcal{D}_f)] \approx p(y_{\text{HF}})$. Our confidence about this prediction can be expressed in form of the variance with respect to the posterior model distribution at a test input $p(f^* | \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f)$, which results in credible intervals on the density function $\mathbb{E}_{f^*} [p(y_{\text{HF}}^* | f^*, \mathcal{D}_f)]$ itself:

$$\mathbb{E}_{f^*} [p(y_{\text{HF}}^* | f^*, \mathcal{D}_f)] = \int_{\Omega_{\text{z}_{\text{LF}}^*}} \int_{\Omega_{f^*}} \underbrace{p(y_{\text{HF}}^* | f^*, \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f) p(f^* | \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f) df^* p(\mathbf{z}_{\text{LF}}^*)}_{\mathbb{E}_{f^*} [p(y_{\text{HF}}^* | f^*, \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f)] \approx p(y_{\text{HF}} | \mathbf{z}_{\text{LF}})} d\mathbf{z}_{\text{LF}}^* \quad (8a)$$

$$\begin{aligned} \mathbb{V}_{f^*} [p(y_{\text{HF}}^* | f^*, \mathcal{D}_f)] &= \mathbb{E}_{f^*} \left[\left(p(y_{\text{HF}}^* | f^*, \mathcal{D}_f) \right)^2 \right] - \left(\mathbb{E}_{f^*} [p(y_{\text{HF}}^* | f^*, \mathcal{D}_f)] \right)^2 \\ &= \int_{\Omega_{\text{z}_{\text{LF}}^*}} \int_{\Omega_{\text{z}_{\text{LF}}^*}'} \mathbb{E}_{f^*} [p(y_{\text{HF}}^* | f^*, \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f) p(y_{\text{HF}}^* | f^*, \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f)] \\ &\quad \cdot p(\mathbf{z}_{\text{LF}}^*) p(\mathbf{z}_{\text{LF}}^*) d\mathbf{z}_{\text{LF}}^* d\mathbf{z}_{\text{LF}}^* - \left(\mathbb{E}_{f^*} [p(y_{\text{HF}}^* | f^*, \mathcal{D}_f)] \right)^2 \end{aligned} \quad (8b)$$

2.3 | Numerical Approximation of Posterior Statistics using Gaussian Processes

In this work, we advocate Gaussian Processes (GPs)²⁹ as a probabilistic regression model in equation (6). GPs are a popular non-parametric Bayesian tool that is well-suited to small data settings. For the inference and learning tasks, we employed *GPpy*³⁰. In all examples analyzed we started with a prior Gaussian Process of the form:

$$p(f) = \mathcal{GP}_f(m(\mathbf{z}_{\text{LF}}), k(\mathbf{z}_{\text{LF}}, \mathbf{z}_{\text{LF}}')), \quad (9a)$$

$$m(\mathbf{z}_{\text{LF}}) = 0, \quad (9b)$$

$$k(\mathbf{z}_{\text{LF}}, \mathbf{z}_{\text{LF}}') = \sigma_0^2 \cdot \exp \left[-\frac{|\mathbf{z}_{\text{LF}} - \mathbf{z}_{\text{LF}}'|^2}{2\ell^2} \right], \quad (9c)$$

where $m(\mathbf{z}_{\text{LF}})$ is the prior mean function and $k(\mathbf{z}_{\text{LF}}, \mathbf{z}_{\text{LF}}')$ is the prior covariance function which we choose to be the squared exponential covariance function with ℓ being the characteristic length scale and σ_0^2 the signal variance. The prior mean function of the process is usually selected to zero or is set equal to y_{LF} , if it is assumed that the LF output reflects y_{HF} in the mean. Following our notation, capital letters, e.g., Γ_i^* , denote the vector of realizations of a random variable, e.g., γ_i and bold capital letters denote matrices, such a matrix holding column-wise realizations of vector valued random variables, e.g. \mathbf{Z}_{LF} , or the covariance matrix \mathbf{K} . The asterisk superscript indicates the large data set derived from the LF Monte-Carlo simulation. Furthermore, we model a Gaussian likelihood of the data with noise level σ_n^2 :

$$p(\mathbf{Y}_{\text{HF}} | F, \mathbf{Z}_{\text{LF}}, \mathcal{D}_f) = \mathcal{N}_{y_{\text{HF}}}(F, \sigma_n^2 I), \quad (10)$$

with F being a particular realization of the GP, evaluated for the feature matrix \mathbf{Z}_{LF} , where each column corresponds to a training point. The posterior GP $p(f|\mathbf{z}_{\text{LF}}, \mathcal{D}_f)$, follows then to²⁹:

$$p(f|\mathcal{D}_f) = \mathcal{GP}_{f|\mathcal{D}_f} \left(m_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}), k_f(\mathbf{z}_{\text{LF}}, \mathbf{z}_{\text{LF}}') \right) \quad (11a)$$

$$m_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}) = m(\mathbf{z}_{\text{LF}}) + \mathbf{k}^T(\mathbf{z}_{\text{LF}}) (\mathbf{K} + \hat{\sigma}_n^2 \mathbf{I})^{-1} (\mathbf{Y}_{\text{HF}}^T - m(\mathbf{z}_{\text{LF}})) \quad (11b)$$

$$k_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}, \mathbf{z}_{\text{LF}}') = k(\mathbf{z}_{\text{LF}}, \mathbf{z}_{\text{LF}}') - k(\mathbf{z}_{\text{LF}}, \mathbf{Z}_{\text{LF}}) [\mathbf{k}(\mathbf{Z}_{\text{LF}}, \mathbf{Z}_{\text{LF}}) + \hat{\sigma}_n^2 \mathbf{I}]^{-1} k(\mathbf{Z}_{\text{LF}}, \mathbf{z}_{\text{LF}}') \quad (11c)$$

Here, $\mathbf{K} = k(\mathbf{Z}_{\text{LF}}, \mathbf{Z}_{\text{LF}})$ and $\mathbf{k} = k(\mathbf{z}_{\text{LF}}, \mathbf{Z}_{\text{LF}})$ are used for compact notation. Point estimates of the hyper-parameters of the model $\theta = \{\ell, \sigma_0^2, \sigma_n^2\}$ are determined by maximizing the marginal likelihood²⁹ and are denoted by $\hat{\theta}$ in the sequel. In addition, we denote the posterior variance of the GP (i.e., the posterior covariance for $\mathbf{z}_{\text{LF}} = \mathbf{z}_{\text{LF}}'$) with $v_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}) = k_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}, \mathbf{z}_{\text{LF}})$. Given a test input \mathbf{z}_{LF}^* , the predictive posterior of the value of the GP at this point, i.e. $f^*(\mathbf{z}_{\text{LF}}^*)$, is given by a normal distribution:

$$p(f^*|\mathbf{z}_{\text{LF}}^*, \mathcal{D}_f) = \mathcal{GP}_{f^*} \left(m_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}^*), k_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}^*, \mathbf{z}_{\text{LF}}^*) \right) = \mathcal{N}_{f^*} \left(m_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}^*), v_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}^*) \right) \quad (12)$$

Furthermore, the Gaussian likelihood in equation (10) implies that the predictive distribution for the corresponding value of the HF model's output, denoted by y_{HF}^* , will be:

$$p(y_{\text{HF}}^*|f^*, \mathbf{z}_{\text{LF}}^*) = \mathcal{N}_{y_{\text{HF}}^*} (f^*, \hat{\sigma}_n^2) \quad (13)$$

The involved densities are summarized in Table 1:

TABLE 1 Applied models for the densities in equations (7) to (8b)

Density	Applied model	Description
$p(f^* \mathbf{z}_{\text{LF}}^*)$	$\mathcal{N}_{f^*} (m(\mathbf{z}_{\text{LF}}^*), v(\mathbf{z}_{\text{LF}}^*))$	Prior GP evaluated at \mathbf{z}_{LF}^*
$p(f^* \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f)$	$\mathcal{N}_{f^*} (m_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}^*), v_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}^*))$	Posterior GP evaluated at \mathbf{z}_{LF}^*
$p(y_{\text{HF}}^* f^*, \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f)$	$\mathcal{N}_{y_{\text{HF}}^*} (f^*, \hat{\sigma}_n^2)$	Likelihood of HF data
$p(\mathbf{z}_{\text{LF}})$	Only samples available	LF distr.
$p(y_{\text{HF}}^* \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f)$		Multi-fidelity conditional
$= \int_{\Omega_{f^*}} p(y_{\text{HF}}^* f^*, \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f) \cdot p(f^* \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f) df^*$	$\mathcal{N}_{y_{\text{HF}}^*} (m_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}^*), v_{\mathcal{D}_f}(\mathbf{z}_{\text{LF}}^*) + \hat{\sigma}_n^2)$	
$p(y_{\text{HF}}^* f^*, \mathcal{D}_f)$		Random process for HF density
$= \int_{\Omega_{\mathbf{z}_{\text{LF}}^*}} p(y_{\text{HF}}^* f^*, \mathbf{z}_{\text{LF}}^*, \mathcal{D}_f) p(\mathbf{z}_{\text{LF}}^*) d\mathbf{z}_{\text{LF}}^*$	-	
$p(y_{\text{HF}}^* \mathcal{D}_f)$		Mean estimate for HF density
$= \mathbb{E}_{f^*} [p(y_{\text{HF}}^* f^*, \mathcal{D}_f)]$	see equation (14)	(in small data regime)

Based on the previous results and given the posterior uncertainty of the GP, we can compute the expected value of the density $\mathbb{E}_{f^*} [p(y_{\text{HF}}^*|f^*, D_f)]$ in equation (8a), by averaging over the posterior of the GP as follows:

$$\begin{aligned} \mathbb{E}_{f^*} [p(y_{\text{HF}}^*|f^*, D_f)] &= \int_{\Omega_{z_{\text{LF}}^*}} \int_{\Omega_{f^*}} \mathcal{N}_{y_{\text{HF}}^*}(f^*, \hat{\sigma}_n^2) \mathcal{N}_{f^*}(\mathbf{m}_{D_f}(\mathbf{z}_{\text{LF}}^*), \mathbf{v}_{D_f}(\mathbf{z}_{\text{LF}}^*)) df^* p(\mathbf{z}_{\text{LF}}^*) d\mathbf{z}_{\text{LF}}^* \\ &= \int_{\Omega_{z_{\text{LF}}^*}} \mathcal{N}_{y_{\text{HF}}^*}(\mathbf{m}_{D_f}(\mathbf{z}_{\text{LF}}^*), \mathbf{v}_{D_f}(\mathbf{z}_{\text{LF}}^*) + \hat{\sigma}_n^2) p(\mathbf{z}_{\text{LF}}^*) d\mathbf{z}_{\text{LF}}^* \\ &\approx \frac{1}{N_{\text{sample}}} \sum_{j=1}^N \mathcal{N}_{y_{\text{HF}}^*}(\mathbf{m}_{D_f}(\mathbf{z}_{\text{LF}j}^*), \mathbf{v}_{D_f}(\mathbf{z}_{\text{LF}j}^*) + \hat{\sigma}_n^2) \end{aligned} \quad (14)$$

A Monte Carlo integration over \mathbf{z}_{LF}^* is used in the last step of equation (14). We note, that this integral, respectively its Monte Carlo approximation, depends on inexpensive LF samples $\mathbf{Z}_{\text{LF}}^* = [\mathbf{z}_{\text{LF}i}^*]$ with $i \in \mathbb{N} : [1, N_{\text{sample}}]$, and for the posterior mean $\mathbf{m}_{D_f}(\mathbf{z}_{\text{LF}}^*)$ and variance $\mathbf{v}_{D_f}(\mathbf{z}_{\text{LF}}^*)$ of the GP, no additional HF runs are needed. Similarly, the (posterior) variance of the sought density of y_{HF} can be computed from (8b) by substituting the density approximations in Table 1 and considering the posterior uncertainty of the GP. Here, we calculate the variance with respect to the GP realizations f^* at \mathbf{z}_{LF}^* and make use of the arithmetic for Gaussian distributions to find a semi-analytic formulation for the variance expression up to the integration over \mathbf{z}_{LF}^* and $\mathbf{z}_{\text{LF}}^{* \prime}$, respectively. Again, the outer integrals over $\Omega_{z_{\text{LF}}^*}$, respectively $\Omega_{z_{\text{LF}}^{* \prime}}$, have to be solved via Monte Carlo integration due to the non-Gaussian distributions $p(\mathbf{z}_{\text{LF}}^*)$, respectively $p(\mathbf{z}_{\text{LF}}^{* \prime})$. The subtrahend $(\mathbb{E}_{f^*} [p(y_{\text{HF}}^*|f^*, D_f)])^2$ can be reused from the previous computation in equation (14). For the subsequent derivation we define the vectors $\mathbf{y}_{\text{HF}}^* = [y_{\text{HF}}^*, y_{\text{HF}}^{* \prime}]^T$ and $\mathbf{f}^* = [f^*, f^{* \prime}]^T$ to denote the support of the multivariate normal distributions, which arise from the multiplication of two univariate normal distributions in the expectation expression $\mathbb{E}_{f^*} [\mathcal{N}_{y_{\text{HF}}^*}(f^*(\mathbf{z}_{\text{LF}}^*), \hat{\sigma}_n^2) \cdot \mathcal{N}_{y_{\text{HF}}^*}(f^*(\mathbf{z}_{\text{LF}}^{* \prime}), \hat{\sigma}_n^2)]$ (see^{29,31} for stochastic calculus):

$$\begin{aligned} \mathbb{V}_{f^*} [p(y_{\text{HF}}^*|f^*, D_f)] &= \int_{\Omega_{z_{\text{LF}}^*}} \int_{\Omega_{z_{\text{LF}}^{* \prime}}} \mathbb{E}_{f^*} [\mathcal{N}_{y_{\text{HF}}^*}(f^*(\mathbf{z}_{\text{LF}}^*), \hat{\sigma}_n^2) \cdot \mathcal{N}_{y_{\text{HF}}^*}(f^*(\mathbf{z}_{\text{LF}}^{* \prime}), \hat{\sigma}_n^2)] p(\mathbf{z}_{\text{LF}}^*) p(\mathbf{z}_{\text{LF}}^{* \prime}) d\mathbf{z}_{\text{LF}}^{* \prime} d\mathbf{z}_{\text{LF}}^* \\ &\quad - (\mathbb{E}_{f^*} [p(y_{\text{HF}}^*|f^*, D_f)])^2 \\ &= \int_{\Omega_{z_{\text{LF}}^*}} \int_{\Omega_{z_{\text{LF}}^{* \prime}}} \int_{\Omega_{f^*}} \mathcal{N}_{\mathbf{y}_{\text{HF}}^*} \left(\begin{bmatrix} f^*(\mathbf{z}_{\text{LF}}^*) \\ f^{* \prime}(\mathbf{z}_{\text{LF}}^{* \prime}) \end{bmatrix}, \begin{bmatrix} \hat{\sigma}_n^2 & 0 \\ 0 & \hat{\sigma}_n^2 \end{bmatrix} \right) \\ &\quad \cdot \mathcal{N}_{f^*} \left(\begin{bmatrix} \mathbf{m}_{D_f}(\mathbf{z}_{\text{LF}}^*) \\ \mathbf{m}_{D_f}(\mathbf{z}_{\text{LF}}^{* \prime}) \end{bmatrix}, \begin{bmatrix} \mathbf{v}_{D_f}(\mathbf{z}_{\text{LF}}^*) & \mathbf{k}_{D_f}(\mathbf{z}_{\text{LF}}^*, \mathbf{z}_{\text{LF}}^{* \prime}) \\ \mathbf{k}_{D_f}(\mathbf{z}_{\text{LF}}^{* \prime}, \mathbf{z}_{\text{LF}}^*) & \mathbf{v}_{D_f}(\mathbf{z}_{\text{LF}}^{* \prime}) \end{bmatrix} \right) df^* \\ &\quad \cdot p(\mathbf{z}_{\text{LF}}^*) p(\mathbf{z}_{\text{LF}}^{* \prime}) d\mathbf{z}_{\text{LF}}^{* \prime} d\mathbf{z}_{\text{LF}}^* - (\mathbb{E}_{f^*} [p(y_{\text{HF}}^*|f^*, D_f)])^2 \\ &= \int_{\Omega_{z_{\text{LF}}^*}} \int_{\Omega_{z_{\text{LF}}^{* \prime}}} \mathcal{N}_{\mathbf{y}_{\text{HF}}^*} \left(\begin{bmatrix} \mathbf{m}_{D_f}(\mathbf{z}_{\text{LF}}^*) \\ \mathbf{m}_{D_f}(\mathbf{z}_{\text{LF}}^{* \prime}) \end{bmatrix}, \begin{bmatrix} \mathbf{v}_{D_f}(\mathbf{z}_{\text{LF}}^*) + \hat{\sigma}_n^2 & \mathbf{k}_{D_f}(\mathbf{z}_{\text{LF}}^*, \mathbf{z}_{\text{LF}}^{* \prime}) \\ \mathbf{k}_{D_f}(\mathbf{z}_{\text{LF}}^{* \prime}, \mathbf{z}_{\text{LF}}^*) & \mathbf{v}_{D_f}(\mathbf{z}_{\text{LF}}^{* \prime}) + \hat{\sigma}_n^2 \end{bmatrix} \right) \\ &\quad \cdot p(\mathbf{z}_{\text{LF}}^*) p(\mathbf{z}_{\text{LF}}^{* \prime}) d\mathbf{z}_{\text{LF}}^{* \prime} d\mathbf{z}_{\text{LF}}^* - (\mathbb{E}_{f^*} [p(y_{\text{HF}}^*|f^*, D_f)])^2 \\ &\approx \frac{1}{N_{\text{sample}}^2} \sum_{i=1}^{N_{\text{sample}}} \sum_{j=1}^{N_{\text{sample}}} \mathcal{N}_{\mathbf{y}_{\text{HF}}^*} \left(\begin{bmatrix} \mathbf{m}_{D_f}(\mathbf{z}_{\text{LF}i}^*) \\ \mathbf{m}_{D_f}(\mathbf{z}_{\text{LF}j}^{* \prime}) \end{bmatrix}, \begin{bmatrix} \mathbf{v}_{D_f}(\mathbf{z}_{\text{LF}i}^*) + \hat{\sigma}_n^2 & \mathbf{k}_{D_f}(\mathbf{z}_{\text{LF}i}^*, \mathbf{z}_{\text{LF}j}^{* \prime}) \\ \mathbf{k}_{D_f}(\mathbf{z}_{\text{LF}j}^{* \prime}, \mathbf{z}_{\text{LF}i}^*) & \mathbf{v}_{D_f}(\mathbf{z}_{\text{LF}j}^{* \prime}) + \hat{\sigma}_n^2 \end{bmatrix} \right) \\ &\quad - (\mathbb{E}_{f^*} [p(y_{\text{HF}}^*|f^*, D_f)])^2 \end{aligned} \quad (15)$$

In the following section, we will discuss the details of the composition of \mathbf{z}_{LF} and additionally provide pseudo-algorithms to summarize all important steps of the Bayesian multi-fidelity approach.

2.4 | Determination of Informative Features $\boldsymbol{\gamma}(\mathbf{x})$ and Optimal Training Set D_f

The vector of lower-fidelity features $\mathbf{z}_{\text{LF}}(\mathbf{x})$ contains apart from low-fidelity model outputs $y_{\text{LF}}(\mathbf{x})$, a few informative feature functions summarized in $\boldsymbol{\gamma}(\mathbf{x})$, so that $\mathbf{z}_{\text{LF}}(\mathbf{x}) = [y_{\text{LF}}(\mathbf{x}), \boldsymbol{\gamma}(\mathbf{x})]^T$. This section is devoted to the definition and computation of

informative features $\boldsymbol{\gamma}(\mathbf{x})$. Former versions of *Bayesian Multi-Fidelity Monte Carlo*^{22,14} have employed $\mathbf{z}_{\text{LF}}(\mathbf{x}) = \mathbf{y}_{\text{LF}}(\mathbf{x})$, so that model inputs \mathbf{x} are entirely filtered through the LF model. The exclusive use of $\mathbf{z}_{\text{LF}}(\mathbf{x}) = \mathbf{y}_{\text{LF}}(\mathbf{x})$ leads, in the general case, to conditional densities that might still be challenging to be sufficiently approximated by probabilistic machine learning models in a *small data* scenario contributing to significant model errors (**error 1**) in Section 2.1) and inaccurate results.

We note that employing the whole input vector $\mathbf{z}_{\text{LF}}(\mathbf{x}) \equiv \mathbf{x}$ as, e.g., in Perdikaris et al.²³ for *nonlinear autoregressive multi-fidelity GP regression (NARGP)*, would only be applicable to low-dimensional \mathbf{x} . In dynamic problems it was shown by Lee et al.³² that it is advantageous to incorporate time t and time derivatives, in the form of time shifts $y_{\text{LF}}(\mathbf{x}, t)$, $y_{\text{LF}}(\mathbf{x}, t + \Delta t)$, of the LF simulation outputs as further features. This idea can directly be integrated in our approach, but without the necessity to treat \mathbf{x} explicitly, by choosing, e.g., $\mathbf{z}_{\text{LF}}(\mathbf{x}, t) = [y_{\text{LF}}(\mathbf{x}, t), t, y_{\text{LF}}(\mathbf{x}, t + \Delta t)]^T$.

The discovery of appropriate $\boldsymbol{\gamma}(\mathbf{x})$ that complement the low-fidelity model outputs $y_{\text{LF}_j}(\mathbf{x})$ can be formulated as an *unsupervised* learning task for which a wealth of linear and non-linear dimensionality reduction techniques have been proposed, in the past^{33,34,35,36,37,38}. We would like to emphasize though, that we do not necessarily seek lower-dimensional representations of \mathbf{x} but rather for the features of \mathbf{x} that are *predictive* of y_{HF} . Even in cases where \mathbf{x} is amenable to (non)linear dimensionality reduction, only a subset of the reduced coordinates may affect the output y_{HF} . Additionally, as $p(y_{\text{HF}}|\mathbf{y}_{\text{LF}})$ is already conditioned on the LF output \mathbf{y}_{LF} , it does not make sense to select $\gamma_i(\mathbf{x})$ that furnish information contained in $\mathbf{y}_{\text{LF}}(\mathbf{x})$ so that a naive supervised learning approach would fail. More specifically, in the context of *BMFMC* advocated in this paper, our goal is to identify a few $\gamma_i(\mathbf{x})$ that, in conjunction with $\mathbf{y}_{\text{LF}}(\mathbf{x})$, will reduce as much as possible the model error source **error 1** in Section 2.1.

In the following, we will not change the overall probabilistic modeling strategy, in our case Gaussian Processes, but rather aim to find a better representation for the multi-fidelity conditional distribution $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}})$ by modifying \mathbf{z}_{LF} through γ_i . As pointed out in **error 1**, the assumption of a specific conditional distribution is not fulfilled in the general case and will hence introduce an error. Additionally, in case of a large variance of the random process $\mathbb{V}_{y_{\text{HF}}}[p(y_{\text{HF}}|\mathbf{z}_{\text{LF}})]$ for y_{HF} , an accurate estimation of this variance in form of the noise level $\hat{\sigma}_n^2$ becomes more challenging. Because feature selection based on $p(y_{\text{HF}}|\mathbf{y}_{\text{LF}})$ is hampered by the low number of HF data points $\{\mathbf{X}, Y_{\text{HF}}\}$, we alternatively choose $\gamma_i(\mathbf{x})$ that explain most of the LF model's output variance $\mathbb{V}_{y_{\text{LF}}}[p(y_{\text{LF}})]$ and argue that such $\gamma_i(\mathbf{x})$ are also important features for model dependency $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}})$, i.e., their explicit treatment in $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}})$ would reduce the variance of the expression the most. By explicit treatment we mean the incorporation of γ_i in \mathbf{z}_{LF} so that the Gaussian Process has γ_i as an additional input variable, rather than implicitly accounting for the effect of γ_i by its contribution to the characteristics of the density $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}})$, where we impose a Gaussian noise assumption. On the other hand, if too many informative features γ_i are added to $p(y_{\text{HF}}|\mathbf{z}_{\text{LF}})$, the resulting space $\Omega_{\mathbf{z}_{\text{LF}}}$ becomes too large to be sufficiently covered by D_f , so that **error 2** will increase due to growing epistemic uncertainty. The effect is schematically depicted in Figure 2. For a larger amount of training data, more informative features γ_i can be added to further decrease the approximation error.

We start by demanding, that $\gamma_i(\mathbf{x})$ is selected from the lower dimensional representation $\hat{\mathbf{x}}$ of the original input \mathbf{x} . Therefore, the model input \mathbf{x} separated into uncorrelated dimensions and correlated dimensions: $\mathbf{x} = [\mathbf{x}_{\text{uncorr}}, \mathbf{x}_{\text{corr},i}]^T$, with $i \in \mathbb{N} : [1, n_{\text{corr}}]$. In this notation $\mathbf{x}_{\text{corr},i}^T$ represent the correlated inputs in form of correlated random variables or random fields. A reduced representation $\hat{\mathbf{x}}$ can be achieved by applying unsupervised dimensionality reduction techniques, in our case a truncated Karhunen-Loève expansion (KLE), on $\mathbf{x}_{\text{corr},i}^T$. We want to emphasize that the dimensionality reduction discussed in this section, was not used to generate the realization of random fields that were deployed in the actual simulations but is rather a post-processing step that is only used in the probabilistic learning approach. Starting point is an eigenproblem for the covariance matrix \mathbf{K}^* of the random field, which is defined by the evaluation of its covariance function for the underlying discretization:

$$\mathbf{K}^* \mathbf{v}_{j,i} = \lambda_{j,i} \mathbf{v}_{j,i} \quad (16)$$

The eigenvectors \mathbf{v}_j define a complete basis, in which we can represent $\mathbf{x}_{\text{corr},i}$ as a linear combination of \mathbf{v}_j , so that the coefficients of the expansion yields:

$$\mathbf{c}_i \approx \mathbf{V}_{\text{trunc},i}^T \cdot \mathbf{x}_{\text{corr},i} - \mathbf{m}_i, \text{ with } i \in \mathbb{N} : [1, n_{\text{corr}}] \quad (17)$$

Here, \mathbf{m}_i is the mean vector of the i -th discretized random field and \mathbf{c}_i a vector of coefficients for the truncated basis $\mathbf{V}_{\text{trunc},i}$. We truncate the series expansion in equation (17) when 95% of the explained variance is reached. The explained variance of the discretized field is defined as:

$$\text{explained variance} := \sum_{i=j}^{n_{\text{trunc}}} \lambda_j / \sum_{j=1}^{d_{\text{corr}}} \lambda_j, \text{ with } d_{\text{corr}} = \dim(\mathbf{x}_{\text{corr},i}) \quad (18)$$

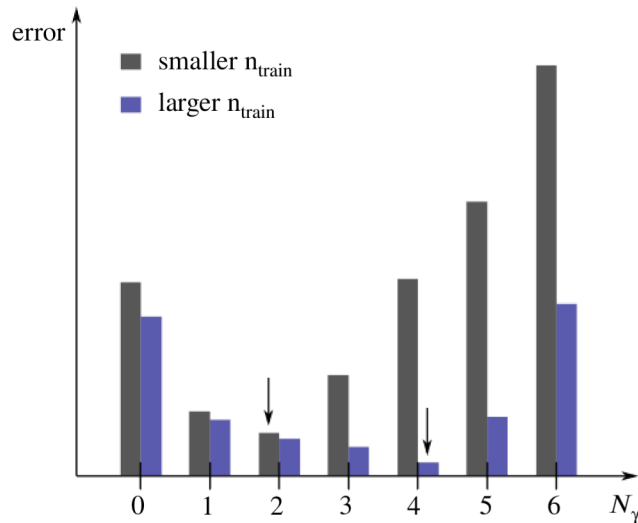


FIGURE 2 Schematic illustration of the error behavior in $p(y_{\text{HF}}|z_{\text{LF}})$ for an increasing number of features, γ_i . The incorporation of features leads at first to a decrease of the modeling **error 1**, before a larger $\Omega_{z_{\text{LF}}}$ leads again to an increase of (**error 2**). For a larger training data size n_{train} the error is in general lower and the minimum error (indicated by arrows) lies at a higher number of features. See Figure 10 for a numerical demonstration.

Afterwards, we propose to use the vector of KLE-coefficients c_i as a low dimensional feature vector for $\mathbf{x}_{\text{corr},i}$. Standardization of each dimension is written in form a standardization operator S . The reduced input vector then follows to:

$$\hat{\mathbf{x}} := S[\mathbf{x}_{\text{uncorr}}, c_i]^T, \text{ with } i \in \mathbb{N} : [1, n_{\text{trunc},i}] \quad (19)$$

Standardization refers to the individual scaling of the dimensions in $\hat{\mathbf{x}}$, so that their underlying input density has zero mean and a standard deviation of one. In a next step, we define a correlation measure \mathbf{r} between the individual dimensions of the reduced input vector $\hat{\mathbf{x}}$ and the LF simulation output y_{LF} , using the projection of the corresponding reduced input matrix $\hat{\mathbf{X}}^*$ on the LF output vector Y_{LF}^* :

$$\mathbf{r} = \left| \hat{\mathbf{X}}^{*T} \cdot Y_{\text{LF}}^* \right| \quad (20)$$

The definition (20) can be understood as the absolute value of the scaled Pearson correlation coefficient, calculated for each dimension of $\hat{\mathbf{x}}$ and the LF output y_{LF} . Input dimensions \hat{x}_j that show high values for r_j are informative about y_{LF} and in approximation about y_{HF} , leading to an efficient reduction of $\mathbb{V}_{y_{\text{HF}}} [p(y_{\text{HF}}|y_{\text{LF}}, \gamma_i)]$. We select the \hat{x}_i that correspond to the i -highest entries r_i in \mathbf{r} , with $i \in \mathbb{N} : [1, n_\gamma]$, as an informative feature γ_i of the input. We found that in the small data regime two additional LF features γ_1 and γ_2 gave the best results. For a higher number of LF features, $\Omega_{z_{\text{LF}}}$ gets too large for the small data set \mathcal{D}_f so that the approximation error increases again.

For an optimal selection of training points \mathcal{D}_f , we propose the following procedure: Given the large LF data set $\mathcal{D}_{\text{LF}}^* := \{\mathbf{X}^*, Y_{\text{LF}}^*\}$, we select a subset of training inputs $\mathbf{X} \subset \mathbf{X}^*$ for which we run HF simulations, according to $Y_{\text{HF}} = y_{\text{HF}}(\mathbf{X})$. The training data set is then defined as $\mathcal{D}_f = \{Y_{\text{HF}}, Y_{\text{LF}}, Z_{\text{LF}}\}$. Important is in the following the definition of the aforementioned subset \mathbf{X} , meaning which \mathbf{x}_i we should select to calculate the training data. Usually, the initial training data selection for surrogate models of computer experiments aims for a space-filling design strategy in the input space, i.e., a Latin Hyper-Cube design to explore the input space efficiently^{39,40,41}. We also decide for a space filling training data set \mathcal{D}_f but pay closer attention to the definition of the term *space filling* in this context as we do not intend to learn a classical deterministic surrogate model but rather want to learn the conditional distribution $p(y_{\text{HF}}|z_{\text{LF}})$ with low error. The noisy relationship between y_{HF} and z_{LF} is caused by neglecting the explicit dependency on \mathbf{x} . To get a good estimate for the noise structure we need to select training points with space filling properties in Ω_x . As the input space is assumed to be large, it is advantageous to focus on the important part of Ω_x , represented by $\gamma^+ = \gamma_i$, with $i \in \mathbb{N} : [1, n_{\gamma^+}]$ and n_{γ^+} being the number of input features used in the extended vector γ^+ . We note that $n_{\gamma^+} > n_\gamma$, such that we demand space filling properties to more dimensions γ_i than actually used to define z_{LF} . This is necessary to yield a training design that is also representative for the noise structure of $p(y_{\text{HF}}|z_{\text{LF}})$. In the extended space Ω_{γ^+} , the sampling data from the LF model $\mathcal{D}_{\text{LF}}^* := \{\mathbf{X}^*, Y_{\text{LF}}^*\}$ corresponds to the LF feature data set Γ^{*+} . Given the training data size n_{train} we then choose a space filling subset $\Gamma^+ \subset \Gamma^{*+}$ from the large sampling data

set of size N_{sample} , with $n_{\text{train}} \ll N_{\text{sample}}$. In the numerical implementation we used *diverse-subset algorithms* based on Wessing and Salomon^{42,43}. Given \mathbf{Z}_{LF}^+ and the corresponding \mathbf{X} , we can now run the HF simulations accordingly to $Y_{\text{HF}} = y_{\text{HF}}(\mathbf{X})$. We found that $n_{\gamma^+} \in \mathbb{N} : [3, 6]$ is a good choice in the small data regime.

Finally, we can summarize the necessary steps for the Bayesian multi-fidelity uncertainty quantification approach in form of the pseudo-code 1. The sub-algorithms 2 and 3 show the details of the determination of \mathbf{z}_{LF} , discussed in this section and the implementation of the posterior statistics from equation (14)-(15).

Algorithm 1 Pseudo-code for Bayesian multi-fidelity UQ (BMFMC)

Require: $p(\mathbf{x})$, $y_{\text{HF}}(\mathbf{x})$, $y_{\text{LF}}(\mathbf{x})$, n_{train} , N_{sample} , $\mathbf{y}_{\text{HF, support}}$

- 1: $\mathbf{X}^* = \text{GENERATE}(p(\mathbf{x}), N_{\text{sample}})$ // Draw N_{sample} Monte-Carlo samples \mathbf{X}^* from $p(\mathbf{x})$
- 2: $Y_{\text{LF}}^* \leftarrow y_{\text{LF}}(\mathbf{X}^*)$ // Run LF model for N_{sample} Monte-Carlo samples
- 3: $p(f|\mathbf{z}_{\text{LF}}) = \text{DESIGNPRIORGp}(Y_{\text{LF}}^*)$
- 4: $\mathcal{D}_f = \text{TRAININGDATA}(\mathbf{X}^*, Y_{\text{LF}}^*, n_{\text{train}}, \mathbf{y}_{\text{HF, support}}, N_{\text{sample}})$ // see algorithm 2
- 5: $\mathcal{GP}_{f|\mathcal{D}_f} \leftarrow \mathcal{GP}_f$ // Train GP model on \mathcal{D}_f

6: **return** $p_{\mathbf{y}_{\text{HF}}, \mathbb{E}^*}$, $p_{\mathbf{y}_{\text{HF}}, \mathbb{V}^*} = \text{POSTERIORSTATISTICS}(\mathcal{GP}_{f|\mathcal{D}_f}, \mathbf{Z}_{\text{LF}}^*, N_{\text{sample}}, \mathbf{y}_{\text{HF, support}})$

Algorithm 2 TrainingData(\mathbf{X}^* , Y_{LF}^* , n_{train} , $\mathbf{y}_{\text{HF, support}}$, N_{sample})

// ———— Unsupervised dim. reduction for input matrix \mathbf{X}^* ———— //

- 1: $[\mathbf{X}_{\text{uncorr}}^*, \mathbf{X}_{\text{corr}, i}^*] \leftarrow \mathbf{X}^*$ // with $i \in \mathbb{N} : [1, n_{\text{corr}}]$, Distinguish: corr. & uncorr. inputs
- 2: **for** i to n_{corr} **do** // Number of random fields: n_{corr}
- 3: $\hat{\mathbf{X}}_i^* = \text{TRUNCATEDKLE}(\mathbf{X}_{\text{corr}, i}^*)$ // Dim. reduction of random fields by truncated KLE
- 4: **end for**
- 5: $\hat{\mathbf{X}}^* \leftarrow [\mathbf{X}_{\text{uncorr}}^*, \hat{\mathbf{X}}_i^*]$ // Construct reduced input matrix
- 6: $\hat{\mathbf{X}}^* = \text{STANDARDIZE}(\hat{\mathbf{X}}^*)$ // Standardize input matrix for line 7

// ———— Define γ^+ (supervised) ———— //

- 7: $\mathbf{r} = \left| \hat{\mathbf{X}}^{*T} \cdot Y_{\text{LF}}^* \right|$ // Calculate corr. coefficients for every dim. in $\hat{\mathbf{x}}$, equation (20)
- 8: **for** i to n_{γ^+} **do** // We found $n_{\gamma^+} \in \mathbb{N} : [3, 6]$ as a good heuristics
- 9: $\text{idx} = \text{RETURNINDEXMAX}(\mathbf{r})$ // Get dimension in $\hat{\mathbf{x}}$ with max. correlation to y_{HF}
- 10: $\Gamma_i^* = \text{SELECTCOLUMN}(\text{idx}, \hat{\mathbf{X}}^*)$ // Select corresponding column in $\hat{\mathbf{X}}^*$
- 11: $\mathbf{r} = \text{SETMAXZERO}(\mathbf{r})$
- 12: **end for**
- 13: $\mathbf{\Gamma}^{*,+} \leftarrow [\Gamma_i^*], i \in \mathbb{N} : [1, n_{\gamma^+}]$ // Construct extended feature space

// ———— Select \mathbf{X} and Y_{HF} ———— //

- 14: $\mathbf{\Gamma}^+ = \text{SELECTDIVERSESUBSET}(\mathbf{\Gamma}^{*,+}, n_{\text{train}})$ // Space filling subset in Ω_{γ^+}
- 15: $\mathbf{X} = \text{GETCORRESPONDINGINPUT}(\mathbf{\Gamma}^+)$
- 16: $Y_{\text{HF}} = y_{\text{HF}}(\mathbf{X})$ // Run HF model for training inputs \mathbf{X}

// ———— Select \mathbf{z}_{LF} and \mathcal{D}_f ———— //

- 17: $\mathbf{z}_{\text{LF}} \leftarrow [\mathbf{y}_{\text{LF}}, \gamma_i]$ // with $i \in \mathbb{N} : [1, n_{\gamma}]$, we recommend $n_{\gamma} = 2$ for $n_{\text{train}} \in \mathbb{N} : [50, 200]$
- 18: $\mathbf{Z}_{\text{LF}} \leftarrow \mathbf{z}_{\text{LF}}$
- 19: $\mathcal{D}_f \leftarrow [\mathbf{Z}_{\text{LF}}, Y_{\text{HF}}]$ // Define training-set for GP of size $n_{\text{train}} \ll N_{\text{sample}}$

20: **return** \mathcal{D}_f

Algorithm 3 PosteriorStatistics($\mathcal{GP}_{f|D_f}, \mathbf{Z}_{\text{LF}}^*, N_{\text{sample}}, \mathbf{y}_{\text{HF,support}}$)

```

// ----- Calculate  $\mathbb{E}_{f^*} [p(y_{\text{HF}}|f^*, \mathbf{D}_f)]$ , equation (14) ----- //
1:  $\mathbf{m}^*, \mathbf{v}^* = \text{EVALUATEGP}(\mathcal{GP}_{f|D_f}, \mathbf{Z}_{\text{LF}}^*)$  // Evaluate the posterior GP at  $\mathbf{Z}_{\text{LF}}^*$ 
2:  $\mathbf{P}_{\text{y}_{\text{HF}}}^* = \text{NORMALPDF}(\mathbf{y}_{\text{HF,support}}, \mathbf{m}^*, \mathbf{v}^*)$  // Matrix with column-wise normal distribution evaluations
3:  $\mathbf{p}_{\text{y}_{\text{HF}}, \mathbb{E}^*} = 1/N_{\text{sample}} \cdot \text{SUMCOLUMNS}(\mathbf{P}_{\text{y}_{\text{HF}}}^*)$  // Sum up normal distributions (same  $\mathbf{y}_{\text{HF,support}}$ )

// ----- Calculate  $\mathbb{V}_{f^*} [p(y_{\text{HF}}|f^*, \mathbf{D}_f)]$ , equation (15) ----- //
4:  $\mathbf{K}_{D_f^*} = \text{POSTERIORCOVARIANCE}(\mathcal{GP}_{f|D_f}, \mathbf{Z}_{\text{LF}}^*)$ 
5:  $j, h = 1$ 
6:  $\mathbf{Y}_{\text{HF}, \mathbb{V}} = [\mathbf{y}_{\text{HF,support}}, \mathbf{y}_{\text{HF,support}}]$  // Evaluate 2-dim. PDF in line 15 for points  $[\mathbf{y}_{\text{HF}, l}, \mathbf{y}_{\text{HF}, l}]^T$ ,
// with  $l \in \mathbb{N} : [1, n_{\text{support}}]$ 

7: for  $\mu_1, v_1$  in  $\mathbf{m}^*, \mathbf{v}^*$  do
8:    $i = 1$ 
9:   for  $\mu_2, v_2$  in  $\mathbf{m}^*, \mathbf{v}^*$  do
10:     $k = \mathbf{K}_{D_f^*}(i, j)$ 
11:     $\hat{\sigma}_n^2 = \text{GETNOISEGP}(\mathcal{GP}_{f|D_f})$ 
12:     $\Sigma = [[v_1 + \hat{\sigma}_n^2, k]^T, [k, v_2 + \hat{\sigma}_n^2]^T]$  // Construct covariance matrix for PDF in line 15
13:     $\mu = [\mu_1, \mu_2]^T$  // Construct mean vector for PDF in line 15
14:     $\mathbf{p}_{\text{y}_{\text{HF}}, \mathbb{V}^*} += \text{NORMALPDF}(\mathbf{Y}_{\text{HF}, \mathbb{V}}, \mu, \Sigma)$  // Evaluate PDF for  $\mathbf{Y}_{\text{HF}, \mathbb{V}}$  and add up results per iteration
15:     $i++ = 1$ 
16:     $h++ = 1$ 
17:   end for
18:    $j++ = 1$ 
19: end for
20:  $\mathbf{p}_{\text{y}_{\text{HF}}, \mathbb{V}^*} = \mathbf{p}_{\text{y}_{\text{HF}}, \mathbb{V}^*} / h - (\mathbf{p}_{\text{y}_{\text{HF}}, \mathbb{E}^*})^2$  // Normalize result and subtract squared mean prediction

21: return  $\mathbf{p}_{\text{y}_{\text{HF}}, \mathbb{E}^*}, \mathbf{p}_{\text{y}_{\text{HF}}, \mathbb{V}^*}$ 

```

3 | SIMPLE GENERATION OF LOW-FIDELITY VARIANTS OF PDE SOLVERS AND ESTIMATION OF COMPUTATIONAL SPEED-UP

This section is devoted to simple and general application-oriented procedures to generate low cost numerical approximators of the original HF model. A HF model is an umbrella term that includes a complex simulation code which captures detailed physical effects, narrow solver tolerances as well as an overall fine numerical discretization of the temporal and spatial domain. LF models could then, for example, be simpler physical models as in²². A physically simpler version of the original problem is however not always easily available so that we investigate theoretical aspects based on pure numerical relaxation of Galerkin methods (Finite Element Method, Discontinuous Galerkin Method).

We first investigate factors that influence the computational cost for approximating the solution of a nonlinear system of partial differential equations based on Galerkin based discretization methods for transient problems with an iterative solution of nonlinear systems of equations within each time step. Afterwards, we motivate a numerical relaxation strategy for the automated generation of LF model versions, based on these considerations. Computational costs can be decomposed in contributions by the present number of degrees of freedom (DoFs), i.e., the number of unknowns arising from the numerical discretization, the necessary number of iterations until convergence, as well as the efficiency of the implementation depending on the computational complexity of a chosen numerical algorithm but also on the optimization level of a specific code. Thus we obtain⁴⁴:

$$\text{cost} \propto \text{DoFs} \cdot \text{time steps} \cdot \text{iterations} \cdot \frac{1}{\text{efficiency of implementation}} \quad (21)$$

We describe the general case of high-order discontinuous Galerkin methods⁴⁵ as these methods contain the widely used finite-element method and finite-volume method as a special case. The spatial dimensionality of the investigated problem is abbreviated by d . The polynomial degree for the Ansatz functions is denoted by k and the measure for the element size by h . The spatial discretization results in N_{ele} elements with $(k+1)^d$ degrees of freedom (and similarly k^d for a continuous finite element space) per element, assuming hexahedral elements and scalar fields. The number of elements is inversely proportional to h to the power of the dimension of the problem: $N_{\text{ele}} \propto \frac{1}{h^d}$. The number of degrees of freedom (DoFs) can then be summarized in equation (22):

$$\text{DoFs} \propto N_{\text{ele}} \cdot (k+1)^d \propto \left(\frac{k+1}{h}\right)^d \quad (22)$$

The cost associated with time stepping is inversely proportional to the time step size (we assume a mean time step for the cost considerations). Besides accuracy demands, the maximal possible time step size is constrained by the stability limits of the deployed solver. As a general discussion of the stability theory for arbitrary solvers is not expedient, we confine the analysis to solvers for transient fluid dynamics applications where the time step is selected according to the CFL condition (*), resulting in⁴⁴

$$\text{time steps} \propto \frac{1}{\Delta t} \stackrel{(*)}{\propto} \underbrace{\frac{k^\gamma}{h}}_{\text{CFL relationship}} \quad \text{with } \gamma \in [1, 2], \quad (23)$$

and argue that the resulting time step size is small enough to ensure a time accurate solution for many problems. Computational costs related to the iterative solution of systems of equations for the unknown degrees of freedom are dependent on solver tolerances ϵ_{solver} , and in general also on the element size h as well as on the polynomial degree k of the Ansatz functions. For *robust solvers* (**), e.g., multigrid, one can assume that the spatial discretization (h, k) does not influence iteration counts. A dependency of computational costs on the solver tolerance according to $-\log \epsilon_{\text{solver}}$ provides a good general model for solvers with optimal complexity, so that we write:

$$\text{iterations} = f(\epsilon_{\text{solver}}, h, k) \stackrel{(**)}{\propto} -\log \epsilon_{\text{solver}} \quad (24)$$

Under *efficiency of implementation* (in DoFs computed per second), we imply the speed at which certain elementary operations of a PDE solver can be performed on a given computer hardware and a given implementation. This factor is summarized in $g(h, k)$ regarding the serial performance of a code, as well as effects of parallel scalability that we summarize in the coefficient of parallel efficiency $\eta_{\text{parallel}}(h, k)$. Furthermore, we introduce the speed-up through floating-point precision $\mathfrak{p} \in \{1, 2\}$, where $\mathfrak{p} = 1$ for (standard) double precision and $\mathfrak{p} = 2$ in case the solver uses single floating-point precision. The factor $g(h, k)$ mainly depends on the implementation variant used for the solver, and we here focus on implementation strategies that have *optimal complexity* (***) w.r.t. the polynomial degree k and mesh size h , so that we can assume the serial performance to be almost independent of these parameters for $k \leq 10$, see⁴⁶, and we further assume *optimal parallel scalability* (***):

$$\text{efficiency of implementation} \propto g(h, k) \cdot \eta_{\text{parallel}}(h, k) \cdot \mathfrak{p} \stackrel{(***)}{\propto} \mathfrak{p} \quad \text{with } \mathfrak{p} = \begin{cases} 1, & \text{for double precision} \\ 2, & \text{for single precision} \end{cases} \quad (25)$$

The speed-up $f_{\text{HF/LF}}$ by a numerical relaxation can then be expressed as:

$$f_{\text{HF/LF}}(k, d, h, \mathfrak{p}) := \frac{\text{costs HF}}{\text{costs LF}} = \frac{\left(\frac{k_0+1}{h_0}\right)^{d_0} \cdot \frac{k_0^\gamma}{h_0} \cdot \frac{1}{\mathfrak{p}_0} \cdot (-\log \epsilon_{\text{solver},0})}{\left(\frac{k+1}{h}\right)^d \cdot \frac{k^\gamma}{h} \cdot \frac{1}{\mathfrak{p}} \cdot (-\log \epsilon_{\text{solver}})} \quad (26)$$

Figure 3 illustrates potential speed-ups due to numerical relaxation of the original problem. Element or mesh coarsening is always shown in combination with a change of the time step size according to the CFL condition. The figure illustrates the effectiveness of higher order methods (change of polynomial degree) in terms of the achieved speed-up factor. A change of k or h changes the number of DoFs according to equation (22). We do not explicitly show theoretical error estimates for the relaxation procedure as they are not of primary importance for the multi-fidelity scheme. It should be noted that the absolute error or deviation between LF and HF model is not relevant for the proposed method, but only the structure of the error itself (noise structure). From a practical perspective, especially the change of the polynomial degree k for the Ansatz function of the Galerkin approximation can lead to large speed-ups even for the same mesh. In contrast, the relaxation of couplings or tolerances has a smaller impact (logarithmic expression). An additional speed-up factor of two can be achieved when we relax the floating point

precision to single-precision. Even though the presented speed-ups are theoretical values one can expect tremendous efficiency gains without the need for a completely new computational model.




FIGURE 3 Example of theoretical speed-ups for a LF model, generated by numerical relaxation. Left: Speed-up over polynomial degree k and mesh-coarsening factor $\frac{h}{h_0}$. The HF reference uses $k_0 = 7$, $\epsilon_{\text{solver},0} = 10^{-6}$; Right: Speed-up over tolerance factor $\frac{\epsilon_{\text{solver}}}{\epsilon_{\text{solver},0}}$ and mesh-coarsening, with HF reference of $k_0 = 2$, $\epsilon_{\text{solver},0} = 10^{-6}$. Along with the spatial coarsening a temporal relaxation is carried out according to the CFL constraint.

The overall cost for the multi-fidelity framework *BMFMC* is then composed by the costs for the sampling on the LF model and costs for the HF simulations in \mathcal{D}_f . To demonstrate the efficiency of *BMFMC*, we want to compare the computational costs of the proposed multi-fidelity approach with the cost associated with a classic Monte-Carlo sampling strategy. We define the following speed-up expression for *BMFMC*:

$$\text{speed-up MF} := \frac{N_{\text{MC,HF}} \cdot \text{costs HF}}{N_{\text{MC,LF}} \cdot \text{costs LF} + n_{\text{train}} \cdot \text{costs HF}} = \frac{N_{\text{MC,HF}} \cdot f_{\text{HF/LF}}}{N_{\text{MC,LF}} + n_{\text{train}} \cdot f_{\text{HF/LF}}} \quad (27)$$

Figure 4 shows the theoretical speed-ups of *BMFMC* for different LF model speed-ups $f_{\text{HF/LF}}$, as well as different training data sizes n_{train} and Monte-Carlo sample sizes N_{MC} . The multi-fidelity approach becomes especially powerful, if a high number of Monte-Carlo evaluations on the HF model would have been necessary to estimate a statistic of interest with high accuracy.

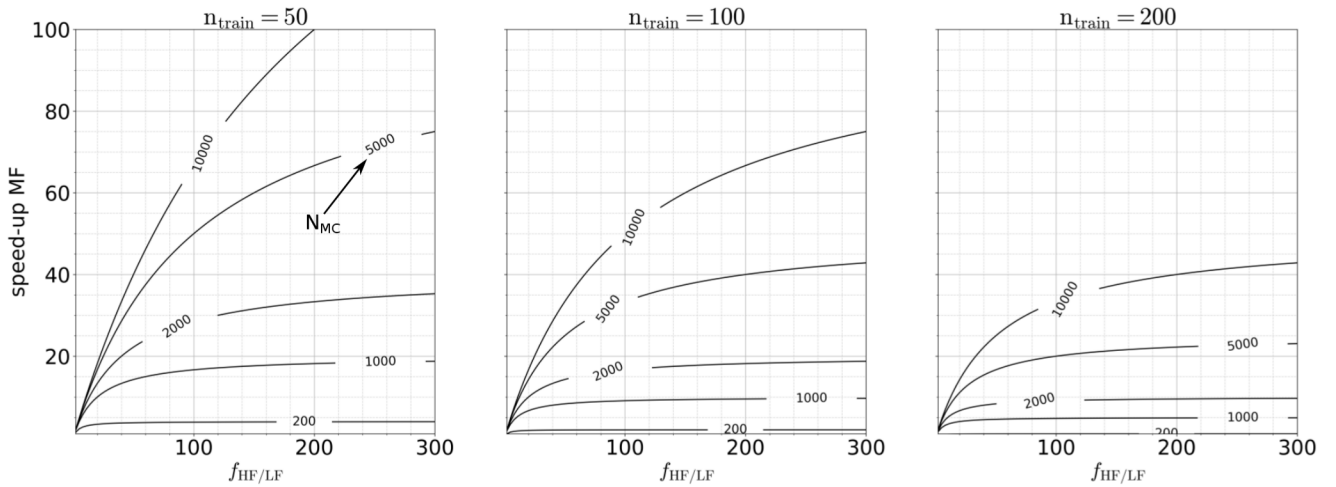


FIGURE 4 Overall speed-up of the proposed multi-fidelity Monte-Carlo (BMFMC) approach for UQ for different training data sizes n_{train} and Monte-Carlo sample N_{MC} sizes as well as different HF/LF model speed-ups.

4 | NUMERICAL DEMONSTRATION FOR CHALLENGING APPLICATIONS

In the following numerical examples, we demonstrate the accuracy and efficiency of the proposed generalized Bayesian multi-fidelity Monte-Carlo framework. The LF models deployed are automatically generated by numerical relaxation of the corresponding HF model, as described in Section 3. The generalized multi-fidelity approach for uncertainty quantification *BMFMC* was implemented in *QUEENS* (Quantification of Uncertainties in Engineering and Science)⁴⁷, a software platform for uncertainty quantification, physics-informed machine learning, Bayesian optimization, inverse problems and simulation analytics. *QUEENS* is capable of interacting with a variety of commercial, open-source and in-house simulation engines and enables the fully automatic set-up of all required simulations on high-performance computing (HPC) clusters, workstations and desktop computers. The first numerical demonstration of a stochastic flow past a cylinder was calculated on a workstation with Intel Core i7-8000K CPUs running at 3.7 GHz and the second demonstration of a stochastic fluid-structure interaction problem was computed on an HPC cluster with Intel Xeon E5-2680v3 "Haswell" CPU running at 2.5 GHz.

4.1 | Stochastic Flow Past a Cylinder: High-Order Discontinuous Galerkin Navier-Stokes Solver

For the first numerical demonstration we investigate uncertainty propagation for a widely used benchmark in computational fluid dynamics: The flow past cylinder test case for incompressible flows, as defined by Schäfer and Turek⁴⁸. Similar setups have also been discussed in the UQ community, see Perdikaris et al.²⁴. The geometry of the two-dimensional domain is a rectangular channel with height $H = 0.41$ and length $L = 2.2$, as depicted in Figure 5. We modify the original benchmark problem at $Re = 100$ to a stochastic flow problem for our investigations on efficient uncertainty propagation: A circular cylinder with uncertain radius $\tilde{R} \sim \mathcal{U}_R(0.035, 0.07)$ is placed in the channel at position $x_c = 0.2$ in streamwise direction. The cylinder's distance to the bottom channel wall is also a univariate random variable $\tilde{y}_c \sim \mathcal{U}_y(0.16, 0.24)$. No-slip boundary conditions are imposed on $\Gamma_{D,0}^F$, defined by the cylinder surface and the channel walls (marked in green in Figure 5). At the outflow boundary Γ_N^F (shown in magenta) a Neumann boundary condition is prescribed as described in⁴⁹. Additionally, the kinematic viscosity $\tilde{\nu}$ of the fluid is uncertain and modeled as a random variable with the univariate distribution $\mathcal{U}_\nu(9.5 \cdot 10^{-4}, 1.5 \cdot 10^{-3})$. Furthermore, we assume a transient and stochastic Dirichlet boundary condition on the inflow section $\Gamma_{D,u}^F$ (shown in blue) in form of a random field in space with a sinusoidal ramp over time:

$$\Gamma_{D,u}^F : \quad \mathbf{u} = \begin{bmatrix} \tilde{u}_x(y, t) \\ 0 \\ 0 \end{bmatrix}, \quad \text{with } \tilde{u}_x(y, t) = \left[\underbrace{U_m \frac{4y(H-y)}{H^2}}_{\text{mean function for } t=T/2} + \underbrace{\mathcal{GP}(0, \tilde{k}(y, y'))}_{\text{Non-stat. random process}} \right] \cdot \underbrace{\sin(\pi t/T)}_{\text{transient ramping}} \quad (28)$$




FIGURE 5 Setup of the stochastic flow past a cylinder problem. Random inputs are written in boxes and have a tilde superscript.

The quantity of interest is the maximal lift coefficient $C_{L,\max}$, due to the lift force L on the cylinder (shown in red) in the y -direction. In the uncertainty propagation problem, we want to infer the distribution $p(C_{L,\max})$ as a stochastic response to the uncertain boundary condition and parameters, whose distribution we abbreviate by $p(\mathbf{x})$. The random process is modeled as a Gaussian Process with a non-stationary kernel function $\tilde{k}(y, y')$, so that 95% of the inflow field's realizations (equal to two times the standard deviation of the process) u_x have less than 25% deviation from the mean $\mu_u(y, t)$. The non-stationary covariance function is formulated as a stationary squared exponential covariance function $k_u(y, y')$ with space dependent signal-variance $(\sigma_u(y))^2$:

$$\tilde{k}(y, y') = \underbrace{(0.125 \cdot \mu_u(y))^2}_{(\sigma_u(y))^2} \cdot \underbrace{\exp\left(-\frac{(y-y')^2}{2\ell^2}\right)}_{k_u(y,y')}. \quad (29)$$

Expression (28) and (29) can be rewritten for easier implementation with available software packages such as *GPY*³⁰:

$$\begin{aligned} \tilde{u}_x(y, t) &\sim \mu_u(y) \cdot \sin(\pi t/T) + 0.125 \cdot \mu_u(y) \cdot \mathcal{GP}(0, k_u(y, y')) \cdot \sin(\pi t/T) \\ &= \mu_u(y) (1 + 0.125 \cdot \mathcal{GP}(0, k_u(y, y'))) \cdot \sin(\pi t/T) \end{aligned} \quad (30)$$

A summary of important properties for the stochastic flow problem is given in Table 2. Discrete realizations of the random

TABLE 2 Properties used in the simulations (random properties are printed bold)

Property	Variable	Value
Channel height	H	0.41
Channel length	L	2.2
Lateral cylinder position	\tilde{y}_c	$\mathcal{U}_y(0.16, 0.24)$
Cylinder radius	\tilde{R}	$\mathcal{U}_R(0.035, 0.07)$
Kinematic viscosity	$\tilde{\nu}$	$\mathcal{U}_\nu(9.5 \cdot 10^{-4}, 1.5 \cdot 10^{-3})$
Inflow BC	$\tilde{u}_x(y, t)$	$\mu_u(y) [1 + 0.125 \cdot \mathcal{GP}(0, k_{u_x}(y, y'))] \cdot \sin(\pi t/T)$
Mean function at $t = T/2$	$\mu_u(y)$	$U_m \frac{4y(H-y)}{H^2}$
Mean max. velocity	U_m	1.5
Correlation length scale	ℓ	$0.08 \cdot H$

inflow field can be computed using standard pseudo-random number generators. We define a vector of points Y_Γ on $\Gamma_{D,u}^P$ on which we evaluate the random inflow BC to yield the velocity vector U_Γ . The Dirichlet boundary condition can then be imposed by a Galerkin projection step.

$$\begin{aligned} \mathbf{u}_\Gamma &= (\mathbf{m} + \mathbf{g}) \cdot \sin(\pi t/T), \text{ with } \mathbf{m} = \mu_u(Y_\Gamma) \text{ and} \\ \mathbf{g} &= \mathbf{L} \cdot \mathbf{r} \sim \mathcal{GP}(\mathbf{0}, k_u(y, y')), \text{ with } \mathbf{r} \sim \mathcal{N}_r(\mathbf{0}, I) \text{ and } \mathbf{L} \cdot \mathbf{L}^T = \mathbf{K}^* \end{aligned} \quad (31)$$

In equation (31) the matrix \mathbf{L} denotes the Cholesky factorization of the covariance matrix $\mathbf{K}^* = k_u(Y_\Gamma, Y_\Gamma')$. The normally distributed vector \mathbf{r} has the dimension of Y_Γ , which in our case was discretized by 200 points. The resulting stochastic dimension (before dimensionality reduction) was $d = 203$. For subsequent parts of our analysis (as put forth in Section 3) we can directly calculate a low dimensional representation $\hat{\mathbf{X}}^*$ of the high-dimensional inputs \mathbf{X}^* by computing a truncated Karhunen-Loève expansion (KLE) (unsupervised dimensionality reduction) of the random field. The random variables \tilde{R} , \tilde{v} and \tilde{y}_c are independent. Hence, their dimension cannot be further reduced. Figure 6 shows on the right side realizations of the random inflow profile for $t = T/2$ (solid lines) along with their truncated KLE approximation of order six (dashed-lines). The bar-chart (left) shows the explained variance over the KLE truncation order. We decided to truncate the extension at order 10 and store the reduced input data of the truncated inflow field and the three random variables as $\hat{\mathbf{X}}^* \in \mathbb{R}^{N_{\text{sample}} \times 13}$, in contrast to the original input data set $\mathbf{X}^* \in \mathbb{R}^{N_{\text{sample}} \times 203}$:

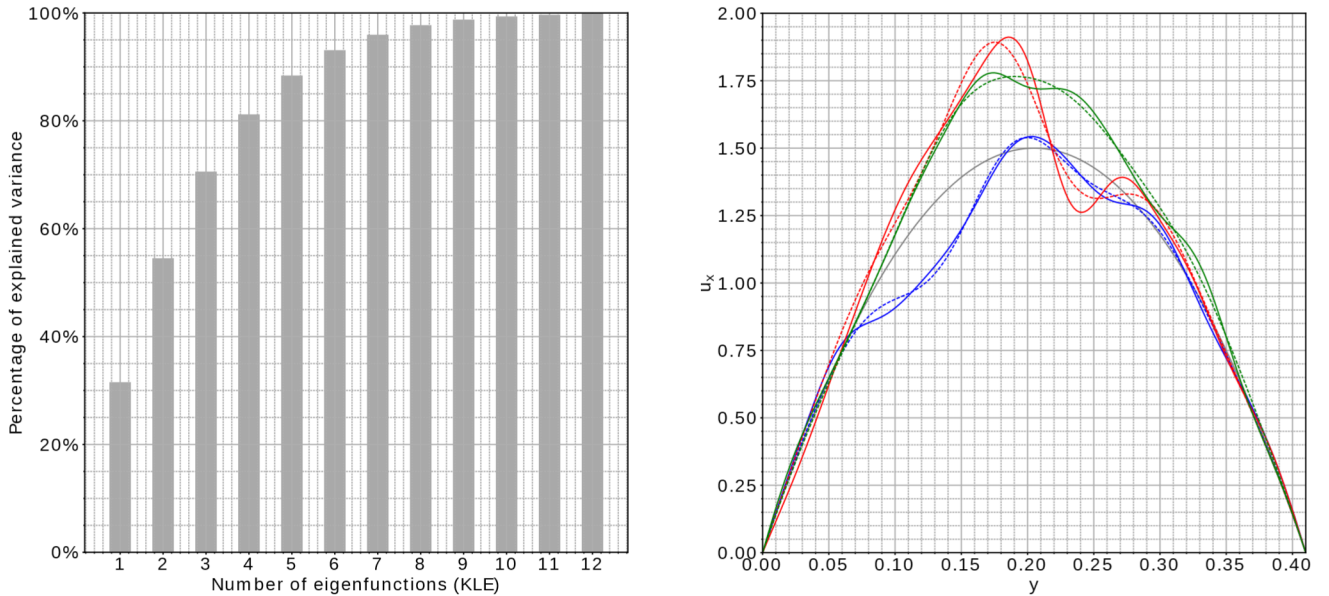


FIGURE 6 Left: Cumulative percentage of explained variance for a Karhunen-Loève expansion of the random inflow field; Right: Example samples (solid line) of the random inflow field $\tilde{u}_x(y)$ for $t = T/2$ along with their Karhunen-Loève approximations (dashed) of order six. The mean function of the random inflow is printed in light grey.

The unsteady problem for a sample realization of the random input $\mathbf{x} \sim p(\mathbf{x})$ is simulated over a time interval of $0 \leq t \leq T = 8$, with a zero velocity field in the domain at initial time $t = 0$. We solve the uncertainty propagation problem using two fidelity levels of a high-order discontinuous Galerkin (L^2 -conforming) discretization developed in^{49,50}. On quadrilateral/hexahedral elements, the solution is approximated by tensor-product Lagrange polynomials of degree $k \geq 2$ for the velocity unknowns, and degree $k_p = k - 1$ for the pressure unknowns for reasons of inf-sup stability. From a practical perspective, the flexibility to vary the polynomial degree k of the shape functions and the mesh resolution h independently, as a means to increase the spatial approximation properties of the discretization, is attractive as one does not have to generate several meshes. We exploit this property for the multi-fidelity approach and define a high-fidelity model version with polynomial degree $k = 6$ and a low-fidelity version of the benchmark using $k = 3$ to resolve the velocity field. For efficient time integration, the method used in the

present work relies on well-known projection methods that segregate the solution of velocity and pressure unknowns, where the second-order accurate dual splitting scheme with an explicit treatment of the convective term is used here.¹

In all simulations, the same parameterized mesh according to Fehn et al.⁴⁹ was used. The simulation domain and the mesh for one sample random input realization $\mathbf{x} \sim p(\mathbf{x})$ is shown in Figure 7 for the HF and the LF model version, respectively. The reduction of the polynomial degree from $k = 6$ to $k = 3$ leads to a speed-up of roughly eight which is in agreement with equation (26). The procedure for the multi-fidelity uncertainty propagation follows algorithm 1. We run simulations for

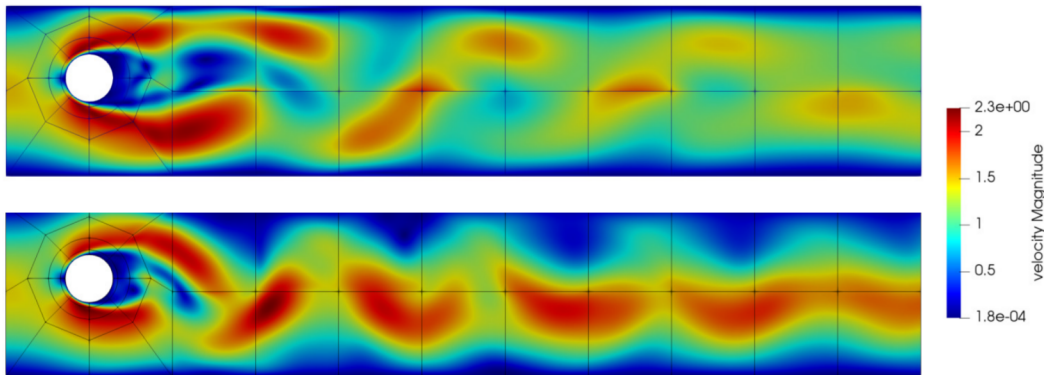


FIGURE 7 Example snapshot of the velocity magnitude for a low-fidelity simulation with $k = 3$ (top) and high-fidelity model simulation with $k = 6$ (bottom). Both simulations used identical inputs \mathbf{x} and are shown for the same simulation time $t = T/2$.

the $N_{\text{sample}} = 10000$ input realizations of $p(\mathbf{x})$ stored in \mathbf{X} , with $\mathbf{x} = [\tilde{u}_x(y), \tilde{v}, \tilde{R}, \tilde{y}_c]^T$ on the LF version of the cylinder flow problem and obtain a vector Y_{LF}^* of according LF model responses for $C_{L,\text{max}}$.

The necessary number of sample points for the LF model is problem-dependent and strongly relies on the goal of the analysis. As we are interested in the entire density, the amount of necessary sample points for an accurate approximation is significantly higher than for point estimates, such as the mean value or the variance. The convergence of the density estimate or statistic of interest can be investigated over an increasing number of sample points. Error *estimates* exist only for point estimators such as Monte-Carlo mean or variance estimators but can be used for an initial orientation of the sample size N_{sample} ⁵¹. The standard error of the Monte-Carlo estimator for the mean yields:

$$\sigma_{\mathbb{E}} \approx \frac{\hat{\sigma}}{\sqrt{N_{\text{sample}}}}, \quad (32)$$

where $\sigma_{\mathbb{E}}$ is the standard deviation of Monte Carlo error, $\hat{\sigma}$ is the estimate of the standard deviation of the QoI and N_{sample} is the number of sample points. We select an initial sample size of $N_{\text{sample}} = 10000$ for the LF model, so that the relative error $\sigma_{\mathbb{E}}/\hat{\sigma}$ in the mean estimate is 1%.

Afterwards, we successively compute features γ_i and choose five features to calculate a $\Omega_{\gamma_i \times y_{\text{LF}}}$ -filling subset $[y_{\text{LF}}(\mathbf{X}), \gamma_i(\mathbf{X})]^T \subset [y_{\text{LF}}(\mathbf{X}^*), \gamma_i(\mathbf{X}^*)]^T$, with $i \in \mathbb{N} : [1, 5]$. We choose a data set of size $n_{\text{train}} = 150$, corresponding to 150 HF model simulations to train the Gaussian Process model. In all problems we investigated, a choice of $n_{\text{train}} \in \mathbb{N} : [50, 200]$ offered a good balance between accuracy and performance. Figure 8 shows the HF and LF model dependency in $\Omega_{y_{\text{HF}} \times y_{\text{LF}}}$ along with the GP-based probabilistic model that would result without γ_i . The Gaussian Process model in $\Omega_{y_{\text{LF}} \times y_{\text{HF}}}$, shown in Figure 8, does not sufficiently explain the complex, non-Gaussian nature of the Monte-Carlo reference, shown by grey dots (normally not available). The introduction of a further dimension γ_i leads to a higher dimensional space in which the data can be better explained by a GP. The transition from $\Omega_{y_{\text{HF}} \times y_{\text{LF}}}$ to $\Omega_{y_{\text{HF}} \times y_{\text{LF}} \times \gamma_1}$ did not require any further HF model evaluation and follows the procedure described in Section 2.4, so that a reduction in the overall approximation error is possible *without further computational efforts*.

The resulting approximation for the HF response $p(y_{\text{HF}} | \mathcal{D}_f)$ is shown in Figure 9a) to d). Figure 9a) shows the best BMFMC prediction for the HF output density, using $n_{\text{train}} = 150$, and two additional features γ_1 and γ_2 . The training points \mathcal{D}_f were selected by choosing a diverse subset in Ω_{γ^+} with $n_{z_{\text{LF}}^+} = 5$. The credible intervals on the densities were computed using

¹To obtain a flow solver that is computationally efficient at high polynomial degrees k , it is essential to implement the method with optimal computational complexity by the use of matrix-free evaluation techniques⁴⁶.

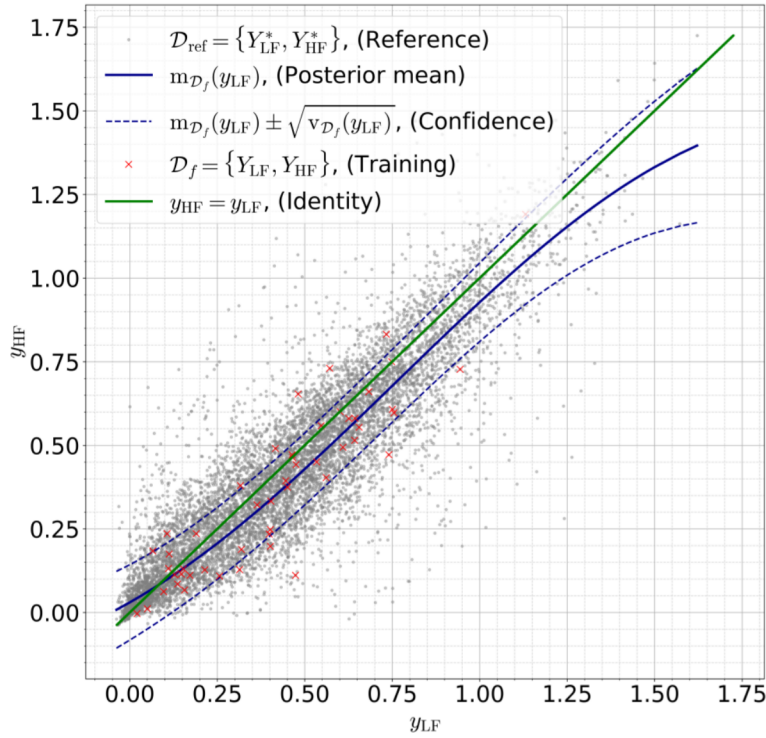


FIGURE 8 LF and HF output tuples of the flow past a cylinder problem. The posterior Gaussian Process for the approximation of $p(f^*|D_f)$ is shown in form of its mean function $m_{D_f}(y_{LF})$ and associated credible intervals. The training data D_f is marked by red crosses.

equation (15), respectively the square root of equation (15) for the standard deviation, and provide an estimate for the uncertainty in the density prediction. The LF density (red line) of the maximum lift coefficient shows a bimodal characteristic that cannot be found in the HF reference density. The BMFMC prediction for $n_{\text{train}} = 150$ without informative LF features in Figure 9b) already resulted in very good predictions. The addition of two informative features in Figure 9a) gave even better predictions for the tails of the distribution and resulted in slightly lower predictive variance (narrower credible intervals) and was in almost perfect agreement with the Monte-Carlo density estimate (dashed black line) that used $N_{\text{sample}} = 10000$ HF evaluations. Figures 9c) and d) used a different strategy to select the training data set D_f : The outcomes of the LF Monte-Carlo simulation Y_{LF} were separated in 25 bins and then an equal amount of training candidates was randomly selected from each bin to define D_f . Even though this strategy covers y_{LF} efficiently, the input \mathbf{x} was not sufficiently covered by the training data, resulting in worse predictions with higher predictive variance. Figure 9d) demonstrates that the predictive variance (credible intervals on the HF density estimates) serves as a good indicator for too small training data size by giving larger credible intervals for a smaller training data size of $n_{\text{train}} = 50$.

To measure the overall accuracy of the predictive HF distribution $\mathbb{E}_{f^*} [p(y_{HF}^*|f^*, D_f)]$, we define an absolute error measure in terms of the Kullback-Leibler divergence (KLD) in equation (33) towards the HF Monte-Carlo density estimate $p(y_{HF})$, which was calculated with a Gaussian kernel density estimation with bandwidth optimization, using $N_{\text{sample}} = 10000$:

$$\epsilon_{\text{abs}} := D_{\text{KL}} \left[p(y_{\text{HF}}) \parallel \mathbb{E}_{f^*} [p(y_{\text{HF}}^*|f^*, D_f)] \right] = \int_{-\infty}^{\infty} p(y_{\text{HF}}) \ln \left(\frac{p(y_{\text{HF}})}{\mathbb{E}_{f^*} [p(y_{\text{HF}}^*|f^*, D_f)]} \right) dy_{\text{HF}} \quad (33)$$

The KLD can be understood as an asymmetric measure of similarity between two probability densities, where two identical distributions would result in a KLD value of zero and a discrepancy between the densities in KLD values greater than zero.

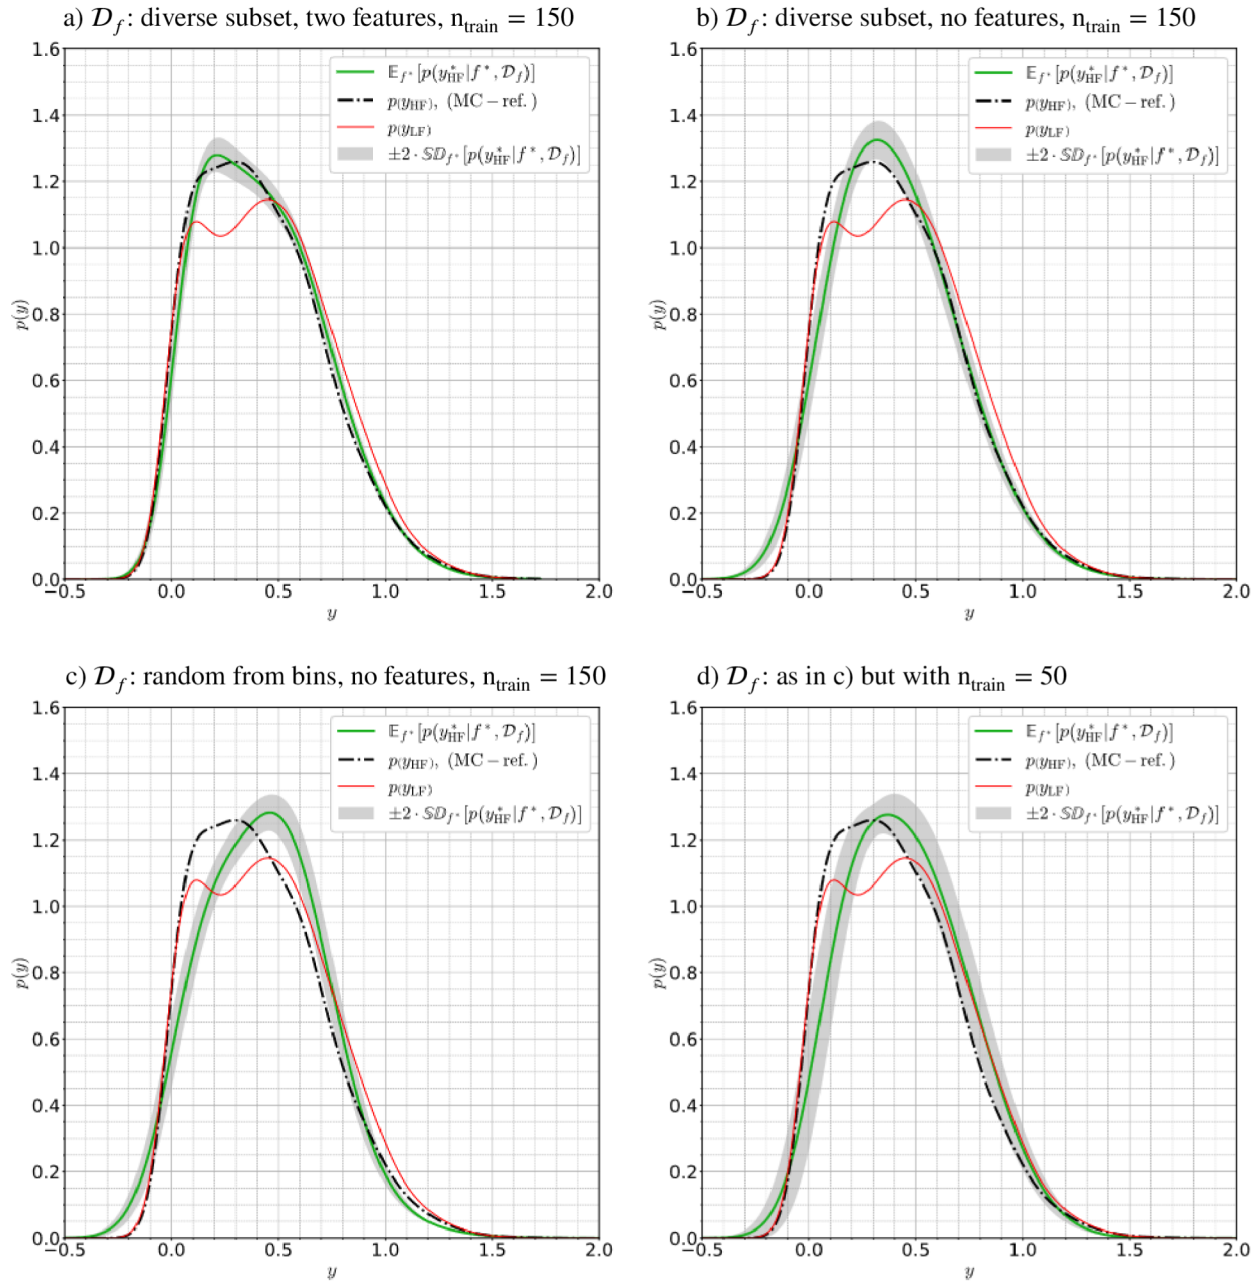


FIGURE 9 Comparison of the output distributions $p(y) = p(C_{L,\max})$ for the maximal lift coefficient in the flow around a cylinder problem. The high-fidelity Monte-Carlo reference $p(y_{\text{HF}})$ is shown as a black dashed line, the low-fidelity solution $p(y_{\text{LF}})$ is given in red color and the BMFMC mean predictions $\mathbb{E}_{f^*}[p(y_{\text{HF}}^*|f^*, \mathcal{D}_f)]$ is shown in green, along with ± 2 standard deviation credible intervals of the prediction, shown in grey.

Figure 10 shows the performance of the Bayesian multi-fidelity approach using the KLD, over an increasing number of features γ_i and two different training sizes n_{train} . To give a reference for the KLDs of the BMFMC solution, we additionally provide the KLD for Monte-Carlo based density estimates using a lower number of sample points (horizontal dashed lines) compared to the Monte-Carlo reference using $N_{\text{sample}} = 10000$.

The KLD for the BMFMC predictions with zero to two additional LF features, which only required 50 HF model simulation, lies far below the KLD that was reached with density estimate, using 5000 HF model evaluations. For comparison we also plotted the reference for a density estimate using only 50 HF simulations, resulting in considerably worse prediction than BMFMC

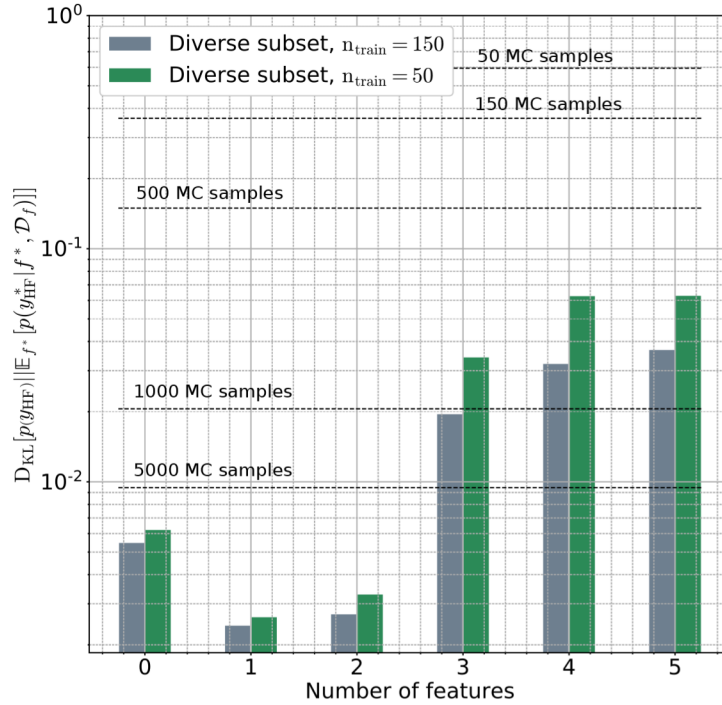


FIGURE 10 Kullback-Leiber divergence between the Monte-Carlo reference solution $p(y_{\text{HF}})$ and the Bayesian prediction $\mathbb{E}_{f^*} [p(y_{\text{HF}}^* | f^*, \mathcal{D}_f)]$ for different number of features γ in \mathbf{z}_{LF} (logarithmic scale). The horizontal dashed lines mark the KL-divergence for one exemplary Monte-Carlo estimate with less points towards the Monte Carlo estimate with $N_{\text{sample}} = 10000$. The Monte-Carlo density estimates yield a significantly higher KLD than the BMFMC estimates.

(please note the logarithmic scale). The best BMFMC predictions were made with only one additional LF feature γ_1 . The use of one or two additional features leads to a significant reduction in the KLD, without the need for any additional computational cost. For the small training data set \mathcal{D}_f with $n_{\text{train}} = 50$ and $n_{\text{train}} = 150$, the introduction of more than two additional features led to a strong increase in the approximation error due to the curse of dimensionality.

The diverse subset strategy for training points selection resulted in considerably better estimates than a random training point selection.

In conclusion, a computational cost comparison of *BMFMC* with a standard Monte-Carlo procedure, as presented in equation (27), resulted in an overall speed-up factor of roughly 7.1, using $N_{\text{MC}} = 10000$, $N_{\text{sample}} = 150$ and $f_{\text{HF/LF}} = 8$, for the problem at hand. In the case of $N_{\text{sample}} = 50$, the speed-up factor was roughly 7.7. The deployed LF model, which was created by changing the degree of the polynomial Ansatz function, is only to be understood as a proof of concept. Much higher speed-ups are possible when further strategies of Section 3 are combined or even a simplified physical model is applied.

4.2 | Stochastic Fluid-Structure Interaction - Bending Wall in a Channel Flow

In the second numerical example, we are interested in uncertainty propagation for a 3D fluid-structure interaction (FSI) problem of a bending wall in a channel flow. The benchmark is motivated by work on monolithic FSI solvers⁵² and is depicted in Figure 11. The fluid domain Ω^F is given by a flow-channel with rectangular cross-section of width $b^F = 1.0$, height $h^F = 0.5$ and length $l^F = 3.0$, while the structure domain Ω^S is represented by an elastic wall of thickness $t^S = 0.05$, width $b^S = 0.6$ and height $h^S = 0.4$. We assume the flexible wall to be clamped to the channel floor at $y = -\frac{h^F}{2}$. The distance between the wall's centerline and the left boundary of the fluid domain is $l_{\text{in}} = 0.5$. The problem is modeled using a hyper-elastic neo-Hookean constitutive law for the structural domain and incompressible Newtonian fluid in the fluid domain. The fluid-structure interaction problem is efficiently solved with a monolithic coupling scheme implemented in our in-house multi-physics finite element code *BACI*⁵³. The interested reader is referred to⁵⁴ for further details on n-field monolithic solvers. In our configuration, the fluid field

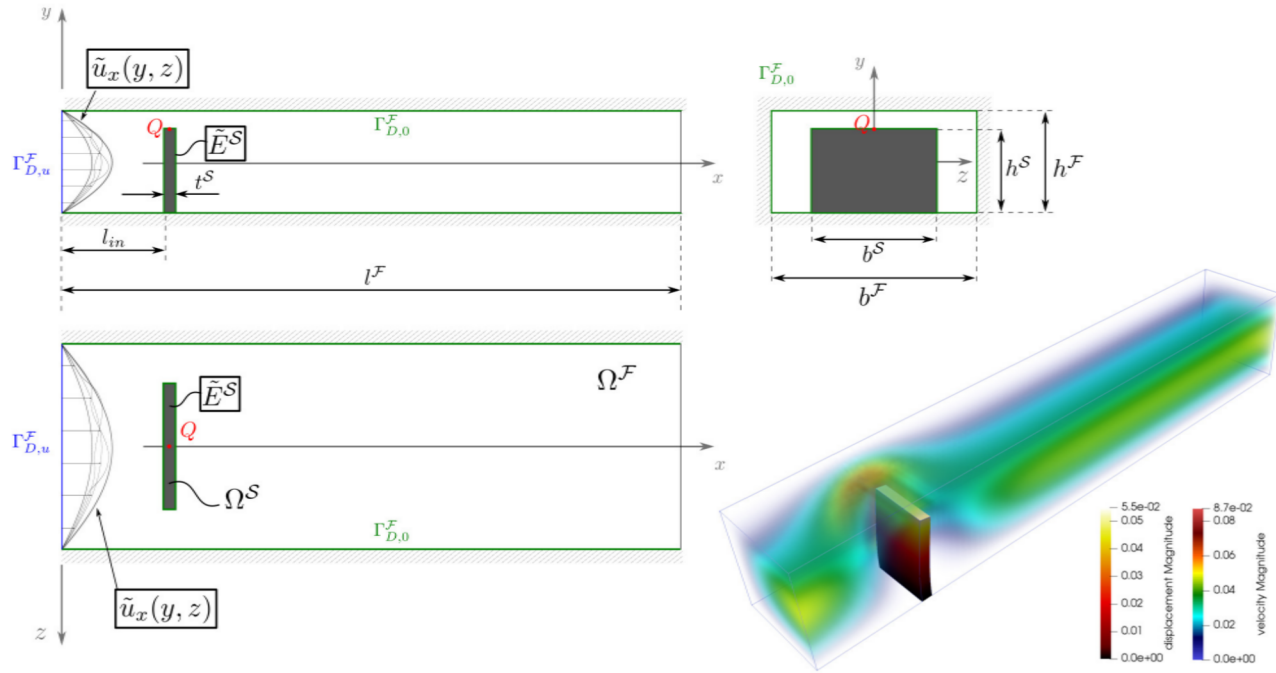


FIGURE 11 Fluid-structure interaction problem of an elastic wall in a channel flow⁵² subject to a random inflow boundary condition and uncertain wall elasticity. Random fields, respectively variables are shown in boxes and have an additional tilde superscript. A no-slip boundary condition is present at $\Gamma_{D,0}^F$. The quantity of interest is the wall deflection in the x-direction at point Q (shown in red). Dirichlet boundaries for the fluid and the structure domain are shown in blue, respectively green.

was chosen as the master field in the dual mortar formulation for the interface coupling⁵⁵. The Dirichlet boundary conditions are formulated on $\Gamma_D^F = \Gamma_{D,u}^F \cup \Gamma_{D,0}^F$ as follows:

$$\begin{aligned} \Gamma_{D,u}^F : \quad \mathbf{u} &= \begin{bmatrix} u_x(y, z) \\ 0 \\ 0 \end{bmatrix} \quad \text{inflow B.C.} \\ \Gamma_{D,0}^F : \quad \mathbf{u} &= \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad \text{no slip B.C.} \end{aligned} \quad (34)$$

An overview of detailed material and fluid properties is provided in Table 3:

The uncertainty propagation problem is then stated as followed: The model is subject to two sources of input uncertainties: A random inflow boundary condition $\tilde{u}_x(y, z)$ (random field with high stochastic dimension) and an uncertain Young's modulus \tilde{E}^S (random variable) for the elastic wall. We are interested in the distribution of the x-displacement of point Q , on top of the elastic wall. To be compliant with our notation in previous sections, we summarize the distribution of the inputs by $p(\mathbf{x})$ and the response distribution for the QoI of the high-fidelity computer model is denoted by $p(y_{\text{HF}})$. Before we investigate the Bayesian multi-fidelity framework for the problem at hand, we discuss the probabilistic models for $p(\mathbf{x})$ in more detail. The **Young's modulus** E^S of the elastic wall is modeled as a random variable with a log-normal distribution $p(E^S) = \mathcal{LN}_{E^S}(\mu_E, \sigma_E^2)$ to constrain realizations to \mathbb{R}^+ . The distribution parameters μ_E and σ_E^2 are chosen so that the Young's modulus mean value is 600 with a standard deviation of 7% from the mean. The **random inflow** $u_x(y, z)$ is realized as a non-stationary Gaussian random field with parabolic mean function on $\Gamma_{D,u}^F$, analogously to the previous numerical example. Again we can factorize the process to:

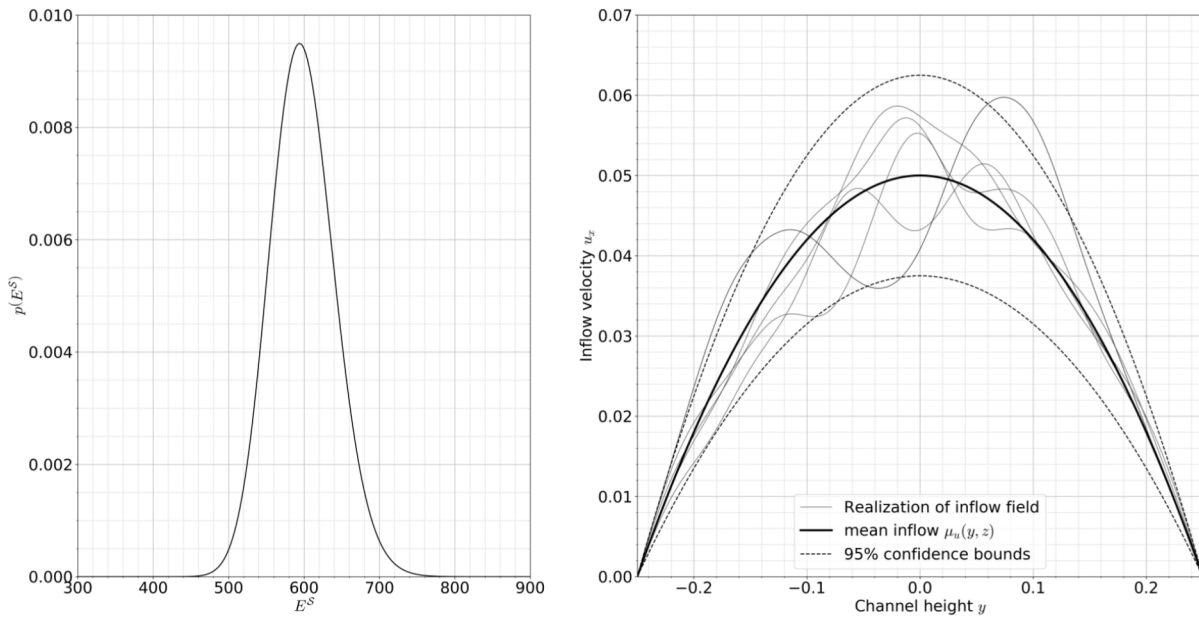
$$u_x(y, z) \sim \mathcal{GP}(\mu_u(y, z), \tilde{k}(y, y')) = \mu_u(y, z) + \mathcal{GP}(0, \tilde{k}(y, y')) \quad (35)$$

The parabolic mean function was taken from the deterministic problem in⁵² and is given in Table 3 along with further properties of the stochastic problem. Densities of the Young's modulus and the resulting random field for the uncertain inflow are visualized

TABLE 3 Material and fluid properties used in the simulations (uncertain properties are printed bold face)

Property	Variable	Value
Poisson ratio	ν^S	0.0
Solid density	ρ^S	1.0
Young's modulus	\tilde{E}^S	$\mathcal{LN}_{E^S}(\mu_E, \sigma_E^2)$
	μ_E	6.392
	σ_E^2	0.00498
Dynamic viscosity	μ^F	0.01
Fluid density	ρ^F	1.0
Inflow Field	$\tilde{u}_x(y, z)$	$\mu_u(\mathbf{x}) \left[1 + 0.125 \cdot \mathcal{GP}(0, k_{u_x}(y, y')) \right]$
Mean function	$\mu_u(y, z)$	$0.05 \cdot \left(1 - \frac{4}{(h^F)^2} y^2 \right) \cdot \left(1 - \frac{4}{(b^F)^2} z^2 \right)$
Stationary covariance function	$k_{u_x}(y, y')$	$\exp \left[\frac{-(y-y')^2}{2\ell} \right]$
Correlation length scale	ℓ	$0.08h^F$

in Figure 12. In the numerical implementation, the realizations of the random inflow were discretized at 200 points so that the stochastic dimension of the problem was $\dim(\mathbf{x}) = 201$. For the stochastic FSI problem, we investigated three different low-

**FIGURE 12** Log-normal density function of the Young's modulus (left) and visualization of a 2D cross-section at $z = 0$ of the 3D-random inflow field (right).

fidelity model versions besides the HF model, regarding their impact on the overall prediction quality. The problem was solved with the continuous finite element method: The high-fidelity model used for the fluid-domain Ω^F 22704 equal-order Hex-8

finite elements with residual based stabilization. The structure domain Ω^S of the HF model used 1536 HEX-8 F-Bar finite elements^{52,56}. The three low-fidelity model variants were constructed using pure numerical relaxation as described in Section 3:

LF 1 The first LF model (LF 1) was designed with a 100 times looser solver tolerances for the fluid-structure coupling and a two times larger time step size resulting in an overall speed-up factor of four.

LF 2 The second LF model (LF 2) was constructed by spatial coarsening to 2838 fluid elements and 192 structural elements, while leaving other numerical settings untouched, leading to a speed-up factor of ten.

LF 3 Finally, the third LF model (LF 3) combined the relaxations of LF 1 and LF 2 and led to a speed-up factor of roughly 28.

The procedure for the Bayesian multi-fidelity uncertainty propagation scheme was then conducted as in the previous example. We first created an input sample matrix $\mathbf{X}^* \sim p(\mathbf{x})$ with $N_{\text{sample}} = 7000$ samples using random number generators for the log-normal distribution, and a Cholesky decomposition from equation (31) to generate sample functions for the random inflow boundary condition. Compared to the previous demonstration, a smaller sample size was chosen, as the smaller variance of the response distribution converged faster. Afterwards, we ran the realizations for \mathbf{x} on all LF models to get the three response vectors \mathbf{Y}_{LF}^* . Features and inputs \mathbf{X} were chosen in the same procedure as in the previous numerical demonstration with the difference that the number of training points was further reduced to $n_{\text{train}} = 50$ HF simulations to demonstrate the capabilities of the multi-fidelity approach even for a very small number of high-fidelity simulations. In case the predictive statistics for the HF output density show too high variance according to equation (15), more training points can be calculated. From a practical perspective we suggest to construct subsets of space filling training designs $\mathbf{X}_1 \subset \mathbf{X}_2 \subset \dots \mathbf{X}_m$ following the procedure in Section 2.4. Afterwards, we can start with the smallest subset (\mathbf{X}_1) and choose larger sets, reusing the previous simulations, in case the predictive variance of the HF density is still too high.

Figure 13a) shows the stochastic dependency between the HF model and LF 1, created by relaxation of the time discretization and coupling settings. The dependency has a non-Gaussian noise structure and strong nonlinearities. The pure spatial relaxation in LF 2 is displayed in Figure 13b). The data points are very close to the identity of $y_{\text{HF}} = y_{\text{LF}}$ (shown in green) along with some scattered points for which the LF simulation deviated from the HF simulation. The combination of relaxations from LF 1 in Figure a) and LF 2 in Figure b) is shown as LF 3 in Figure c). Despite having the highest numerical relaxation and therefore the highest computational speed-up, the LF 3 model results in a less noisy model dependency structure in $\Omega_{y_{\text{HF}} \times y_{\text{LF}}}$ when compared to the LF 1 model. A possible explanation for this effect could be the smoothing property of the coarser mesh which might damp the model discrepancy over the input space. A detailed investigation of the effect is however outside the scope of this paper.

Figure 14 presents the resulting density predictions for the HF output along with the MC reference $p(y_{\text{HF}})_{\text{MC}}$ and the output densities for LF 1 in Figure a) and LF 3 in Figure b), respectively. Even though the LF 1 model led to a very noisy and non-Gaussian model dependency with the HF model, as shown in Figure 13a), the multi-fidelity approach BMFMC was able to predict the HF output density nearly perfectly with as little as 50 HF simulation runs. Additionally, we present the prediction quality of BMFMC without the use of additional informative features (green dashed line) which results in considerably worse predictions due to a high modeling error when assuming a Gaussian noise structure between HF and LF 1 in $\Omega_{y_{\text{LF}} \times y_{\text{HF}}}$. In fact, without the use of γ the modeling error of the GP is so high that the HF reference density did not lie within the predicted credible intervals, even for an increase in training data as demonstrated in Figure 13c). The introduction of only one additional feature γ_1 removed this problem and the reference solution is always within the credible intervals. Figure 13 shows the superior result for two additional features γ_1 and γ_2 that also gave tighter credible intervals than the prediction with only γ_1 . In the case of the LF 3 model, which combined spatial and temporal relaxation, BMFMC predicted the reference density nearly perfectly with or without additional features.

Finally, we investigate the computational speed-ups reached with *BMFMC* in comparison to a standard Monte Carlo strategy. We use the presented speed-up definition in equation (26) and (27) to calculate the speed-up factor. In our simulations we set $N_{\text{HF}} = N_{\text{LF}}$ and the resulting speed-up factors are summarized in Table 4. Only based on pure numerical relaxation, *BMFMC* performed roughly 23 times faster than the Monte Carlo benchmark while reaching the same accuracy. We want to emphasize that we did not even fully exhaust the potential in the numerical relaxation, as i.e., the floating point precision in the simulation was kept untouched. Furthermore, a huge speed-up potential is still available in terms of a simplified physics for the model so that the discussed problems do not represent the full potential of *BMFMC*.

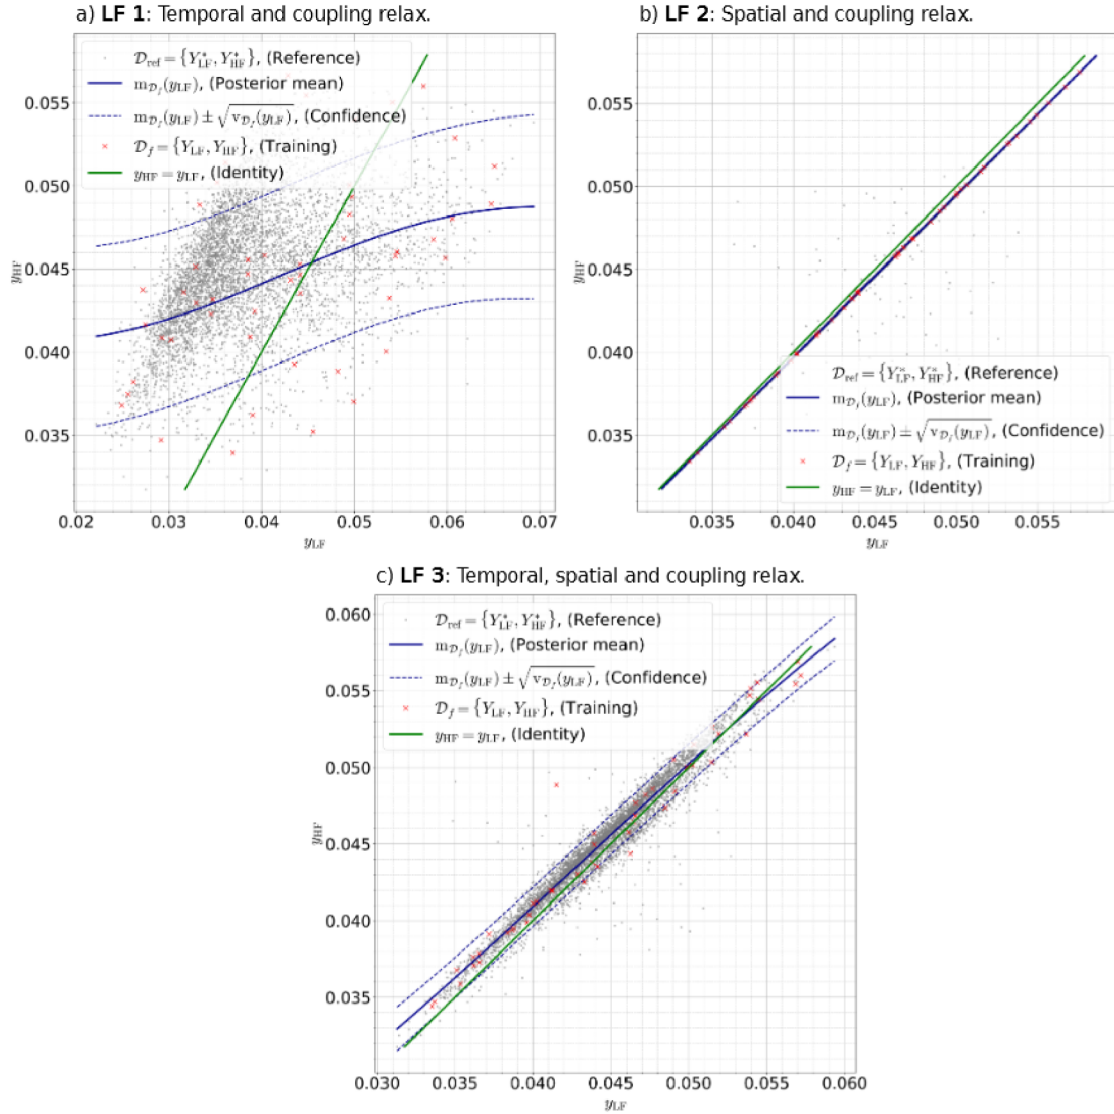


FIGURE 13 Stochastic fluid-structure interaction problem: HF and LF model outputs for the wall deflection, demonstrated for the three low-fidelity models LF 1, LF 2 and LF 3. The data of the Monte-Carlo reference (usually not available) is shown as grey dots. Training points ($n_{\text{train}} = 50$) for the Gaussian Process are shown by red crosses. The posterior mean function $m_{D_j}(y_{\text{LF}})$ and $\pm 1 - \sigma$ standard deviation of the trained Gaussian Process are shown in blue.

TABLE 4 Comparison of computational speed-ups for the generalized *BMFMC* approach with a standard Monte-Carlo approach for uncertainty quantification

LF model	$f_{\text{HF/LF}}$	N_{MC}	n_{train}	speed-up MF
LF 1	4.5	7000	50	4.4
LF 2	10	7000	50	9.3
LF 3	28	7000	50	23.3


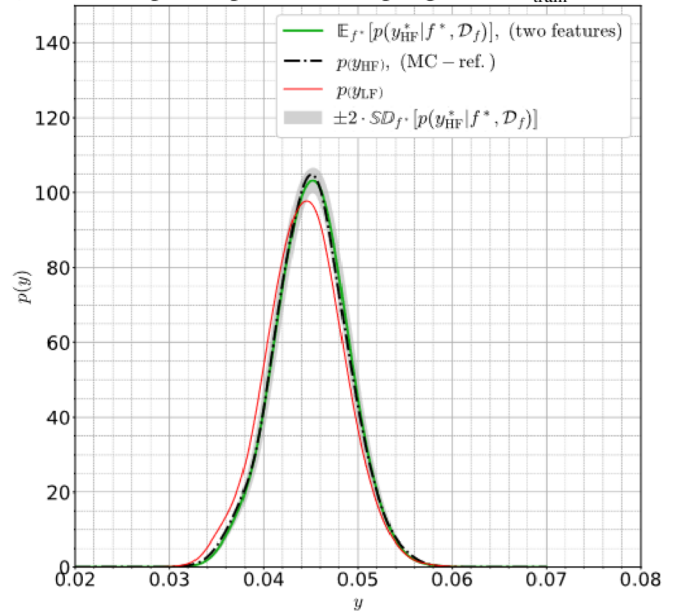


a) **LF 1:** Temporal and coupling relax., $n_{\text{train}} = 50$ b) **LF 3:** Temporal, spatial and coupling relax., $n_{\text{train}} = 50$ c) **LF 1:** Temporal and coupling relax., $n_{\text{train}} = 150$ d) **LF 3:** Temporal, spatial and coupling relax., $n_{\text{train}} = 150$ 

FIGURE 14 Comparison of the output densities $p(y)$ for the bending wall in a channel flow problem. The quantity of interest is the x-deflection point Q . The high-fidelity Monte-Carlo reference density $p(y_{\text{HF}})$ is shown as a black dashed line, the low-fidelity density $p(y_{\text{LF}})$ is given in red and the Bayesian predictions $\mathbb{E}_{f^*}[p(y_{\text{HF}}^*|f^*, D_f)]$, in green along with their credible intervals, shown in grey. Figure a) shows the output densities when LF 1 was deployed in the multi-fidelity approach and Figure b) when LF 3 was deployed, each for $n_{\text{train}} = 50$. Figure c) and d) show the predictions for $n_{\text{train}} = 150$.

5 | CONCLUSION

In this contribution, we introduced a generalized formulation for Bayesian multi-fidelity forward and backward uncertainty propagation. In the case of uncertainty quantification, we call the approach Bayesian multi-fidelity Monte-Carlo or short *BMFMC*. We presented a formulation for the generalized approach that can be used with any probabilistic machine learning model and gave the specific numerical implementation under the use of Gaussian Processes. We demonstrated that the generalized *BMFMC* formulation contains other state-of-the-art methods for uncertainty quantification as special cases, but shows superior performance, especially for computationally expensive UQ problems with high stochastic dimension.

Given a very limited amount of HF training data D_f and a potentially very high stochastic input dimension of \mathbf{x} , both classical and multi-fidelity surrogate based approaches for UQ, that consider \mathbf{x} explicitly, show serious accuracy problems. The reason is an insufficient coverage of $\Omega_{\mathbf{z}_{\text{LF}} \times \mathbf{y}_{\text{HF}}}$ by D_f , with $\mathbf{z}_{\text{LF}} = [y_{\text{LF}}, \mathbf{x}]^T$, in this scenario. For the other extreme, if we only consider $\mathbf{z}_{\text{LF}} = y_{\text{LF}}$, a sufficient space coverage of $\Omega_{\mathbf{y}_{\text{LF}}}$ by training data D_f is possible, but the dependency between the HF and the LF model might still be too complex in $\Omega_{\mathbf{y}_{\text{LF}} \times \mathbf{y}_{\text{HF}}}$. An approximation of the density $p(y_{\text{HF}} | y_{\text{LF}})$ for the small data case is hence prone to modeling errors. Instead, we showed that additional features of the input $\boldsymbol{\gamma}$ can be learned at no extra cost, so that $\mathbf{z}_{\text{LF}} = [y_{\text{LF}}, \boldsymbol{\gamma}]^T$ defines a space $\Omega_{\mathbf{z}_{\text{LF}} \times \mathbf{y}_{\text{HF}}}$, in which the statistical dependence with the HF model response can be learned in a *small data* regime and with a low approximation error. We furthermore demonstrated that our Bayesian approach provides credible intervals for the density estimate itself so that it is possible to assess the quality of the predictions.

We aimed at applicability of *BMFMC* towards complex physical models with actual engineering relevance and demonstrated the capabilities and generality of the method on two challenging stochastic fluid-, respectively, fluid-structure interaction problems with high stochastic input dimensions. We compared the performance of *BMFMC* with a Monte Carlo sampling strategy, which is currently the only reliable alternative for high-dimensional stochastic problems, and demonstrated significant speed-up factors of over 23 when *BMFMC* was used. The speed-ups were achieved by only using simple numerical relaxation of the original problem. We want to emphasize that higher performance gains are possible by, e.g., using simplified physics in the LF model, a simpler geometrical model or combinations of the numerical relaxations discussed in this paper.

In future work the presented method will be extended to the full solution field of an HF model, using more flexible probabilistic modeling approaches. Furthermore, an in-depth investigation of the Bayesian multi-fidelity approach for backward uncertainty propagation should yield significant computational speed-ups. Another interesting application of the method are multi-scale problems where the different physical scales can be interpreted as fidelity levels, keeping the central idea of learning $p(y_{\text{HF}} | \mathbf{z}_{\text{LF}})$ identical to the discussed UQ problems.

ACKNOWLEDGMENTS

The research presented in this paper was partly funded by the German Research Foundation (DFG) under the priority program SPP 1886 "Polymorphic uncertainty modeling for the numerical design of structures" with additional contributions that were funded by the DFG priority program "Software for Exascale Computing" (SPPEXA). The software QUEENS was provided by the courtesy of AdCo Engineering^{GW} GmbH, which is gratefully acknowledged. We finally want to thank Peter Münch for his support on the side of the DG-Navier-Stokes solver, Sebastian Brandstätter for his helpful contributions in the software framework *QUEENS* and Gil Robalo Rei for his support in model preparation.

References

1. Biehler J, Mäck M, Nitzler J, Hanss M, Koutsourelakis PS, Wall WA. Multi-Fidelity Approaches for Uncertainty Quantification. *Surveys for Applied Mathematics and Mechanics (GAMM Mitteilungen)* 2019. doi: <https://doi.org/10.1002/gamm.201900008>
2. Peherstorfer B, Willcox K, Gunzburger M. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review* 2018; 60(3): 550–591. doi: <http://dx.doi.org/10.1137/16M1082469>
3. Heinrich S. Multilevel Monte Carlo methods. In: *Large-Scale Scientific Computing*. Springer, Berlin, Heidelberg. ; 2001: 58–67

4. Giles MB, Nagapetyan T, Ritter K. Adaptive Multilevel Monte Carlo Approximation of Distribution Functions. *arXiv preprint arXiv:1706.06869* 2017.
5. Giles MB. Multilevel Monte Carlo Path Simulation. *Operations Research* 2008; 56(3): 607–617. doi: <http://dx.doi.org/10.1287/opre.1070.0496>
6. Giles MB. Multilevel Monte Carlo Methods. *Acta Numerica* 2015; 24: 259–328. doi: <http://dx.doi.org/10.1017/S096249291500001X>
7. Bierig C, Chernov A. Approximation of probability density functions by the multilevel Monte Carlo maximum entropy method. *Journal of Computational Physics* 2016; 314: 661–681. doi: <https://doi.org/10.1016/j.jcp.2016.03.027>
8. Peherstorfer B, Cui T, Marzouk Y, Willcox K. Multifidelity importance sampling. *Computer Methods in Applied Mechanics and Engineering* 2016; 300: 490–509. doi: <http://dx.doi.org/10.1016/j.cma.2015.12.002>
9. Cui T, Marzouk YM, Willcox KE. Data-driven model reduction for the Bayesian solution of inverse problems. *International Journal for Numerical Methods in Engineering* 2015; 102(5): 966–990. doi: <http://dx.doi.org/10.1002/nme.4748>
10. Li J, Marzouk YM. Adaptive construction of surrogates for the Bayesian solution of inverse problems. *SIAM Journal on Scientific Computing* 2014; 36(3): A1163–A1186. doi: <http://dx.doi.org/10.1137/130938189>
11. Christen JA, Fox C. Markov chain Monte Carlo using an approximation. *Journal of Computational and Graphical Statistics* 2005; 14(4): 795–810. doi: <http://dx.doi.org/10.1198/106186005X76983>
12. Tierney L, Mira A. Some adaptive Monte Carlo methods for Bayesian inference. *Statistics in Medicine* 1999; 18(17-18): 2507–2515. doi: [http://dx.doi.org/10.1002/\(SICI\)1097-0258\(19990915/30\)18:17/18<2507::AID-SIM272>3.0.CO;2-J](http://dx.doi.org/10.1002/(SICI)1097-0258(19990915/30)18:17/18<2507::AID-SIM272>3.0.CO;2-J)
13. Fox C, Nicholls G. Sampling conductivity images via MCMC. *The Art and Science of Bayesian Image Analysis* 1997: 91–100.
14. Koutsourelakis PS. Accurate uncertainty quantification using inaccurate computational models. *SIAM Journal on Scientific Computing* 2009; 31(5): 3274–3300. doi: <http://dx.doi.org/10.1137/080733565>
15. Bilionis I, Zabaras N. Solution of inverse problems with limited forward solver evaluations: A Bayesian perspective. *Inverse Problems* 2013; 30(1): 015004. doi: <http://dx.doi.org/10.1088/0266-5611/30/1/015004>
16. Le Gratiet L, Garnier J. Recursive Co-Kriging model for design of computer experiments with multiple levels of fidelity. *International Journal for Uncertainty Quantification* 2014; 4(5). doi: <http://dx.doi.org/10.1615/Int.J.UncertaintyQuantification.2014006914>
17. Le Gratiet L. Bayesian analysis of hierarchical multifidelity codes. *SIAM Journal on Uncertainty Quantification* 2013; 1(1): 244–269. doi: <http://dx.doi.org/10.1137/120884122>
18. Quaglino A, Pezzuto S, Koutsourelakis PS, Auricchio A, Krause R. Fast uncertainty quantification of activation sequences in patient-specific cardiac electrophysiology meeting clinical time constraints. *International Journal for Numerical Methods in Biomedical Engineering* 2018; 34(7): e2985. doi: 10.1002/cnm.2985
19. Raissi M, Perdikaris P, Karniadakis GE. Inferring solutions of differential equations using noisy multi-fidelity data. *Journal of Computational Physics* 2017; 335: 736–746. doi: <http://dx.doi.org/10.1016/j.jcp.2017.01.060>
20. Takhtaganov T, Müller J. Adaptive Gaussian process surrogates for Bayesian inference. *arXiv preprint arXiv:1809.10784* 2018.
21. Perdikaris P, Karniadakis GE. Model inversion via multi-fidelity Bayesian optimization: a new paradigm for parameter estimation in haemodynamics, and beyond. *Journal of The Royal Society Interface* 2016; 13(118): 20151107. doi: <http://dx.doi.org/10.1098/rsif.2015.1107>

22. Biehler J, Gee MW, Wall WA. Towards efficient uncertainty quantification in complex and large-scale biomechanical problems based on a Bayesian multi-fidelity scheme. *Biomechanics and Modeling in Mechanobiology* 2015; 14(3): 489–513. doi: <http://dx.doi.org/10.1007/s10237-014-0618-0>
23. Perdikaris P, Raissi M, Damianou A, Lawrence N, Karniadakis GE. Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 2017; 473(2198): 20160751. doi: <http://dx.doi.org/10.1098/rspa.2016.0751>
24. Perdikaris P, Venturi D, Royset JO, Karniadakis GE. Multi-fidelity modelling via recursive Co-Kriging and Gaussian–Markov random fields. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 2015; 471(2179): 20150018. doi: <http://dx.doi.org/10.1098/rspa.2015.0018>
25. Grigo C, Koutsourelakis PS. Bayesian Model and Dimension Reduction for Uncertainty Propagation: Applications in Random Media. *SIAM/ASA Journal on Uncertainty Quantification* 2019; 7(1): 292–323. doi: <http://dx.doi.org/10.1137/17M1155867>
26. Kennedy MC, O’Hagan A. Predicting the output from a complex computer code when fast approximations are available. *Biometrika* 2000; 87(1): 1–13. doi: <http://dx.doi.org/10.1093/biomet/87.1.1>
27. Rasmussen CE, Ghahramani Z. Bayesian Monte Carlo. *Advances in Neural Information Processing Systems* 2003: 505–512.
28. Murphy KP. *Machine Learning: A Probabilistic Perspective*. The MIT Press . 2012.
29. Williams CK, Rasmussen CE. *Gaussian processes for machine learning*. 2. MIT press Cambridge, MA . 2006.
30. GPy . GPy: A Gaussian process framework in python. <http://github.com/SheffieldML/GPy>; since 2012.
31. Bertsekas DP, Tsitsiklis JN. *Introduction to probability*. 1. Athena Scientific Belmont, MA . 2002.
32. Lee S, Dietrich F, Karniadakis GE, Kevrekidis IG. Linking Gaussian Process regression with data-driven manifold embeddings for nonlinear data fusion. *arXiv preprint arXiv:1812.06467* 2018.
33. Roweis ST, Saul LK. Nonlinear dimensionality reduction by locally linear embedding. *Science* 2000; 290(5500): 2323–2326. doi: <http://dx.doi.org/10.1126/science.290.5500.2323>
34. Van Der Maaten L, Postma E, Herik V. dJ. Dimensionality reduction: a comparative. *Journal of Machine Learning Research* 2009; 10(66-71): 13.
35. Silva VD, Tenenbaum JB. Global versus local methods in nonlinear dimensionality reduction. In: *Advances in Neural Information Processing Systems*. The MIT Press. ; 2003: 721–728.
36. Tipping ME, Bishop CM. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 1999; 61(3): 611–622. doi: <http://dx.doi.org/10.1111/1467-9868.00196>
37. Wold S, Esbensen K, Geladi P. Principal component analysis. *Chemometrics and Intelligent Laboratory Systems* 1987; 2(1-3): 37–52. doi: [http://dx.doi.org/10.1016/0169-7439\(87\)80084-9](http://dx.doi.org/10.1016/0169-7439(87)80084-9)
38. Höskuldsson A. PLS regression methods. *Journal of Chemometrics* 1988; 2(3): 211–228.
39. Kang F, Han S, Salgado R, Li J. System probabilistic stability analysis of soil slopes using Gaussian process regression with Latin hypercube sampling. *Computers and Geotechnics* 2015; 63: 13–25. doi: <https://doi.org/10.1016/j.compgeo.2014.08.010>
40. Jones B, Johnson RT. Design and analysis for the Gaussian process model. *Quality and Reliability Engineering International* 2009; 25(5): 515–524. doi: <https://doi.org/10.1002/qre.1044>
41. Ranjan P, Haynes R, Karsten R. A computationally stable approach to Gaussian process interpolation of deterministic computer simulation data. *Technometrics* 2011; 53(4): 366–378. doi: <https://doi.org/10.1198/tech.2011.09141>
42. Wessing S. *Two-stage methods for multimodal optimization*. PhD thesis. Universitätsbibliothek Dortmund, 2015.

43. Salomon S, Avigad G, Goldvard A, Schütze O. PSA—a new scalable space partition based selection algorithm for MOEAs. In: *EVOLVE-A Bridge between Probability, Set Oriented Numerics, and Evolutionary Computation II*. Springer. 2013 (pp. 137–151)
44. Fehn N, Wall WA, Kronbichler M. Efficiency of high-performance discontinuous Galerkin spectral element methods for under-resolved turbulent incompressible flows. *International Journal for Numerical Methods in Fluids* 2018; 88(1): 32-54. doi: 10.1002/flid.4511
45. Hesthaven JS, Warburton T. *Nodal discontinuous Galerkin methods: algorithms, analysis, and applications*. Springer . 2007
46. Kronbichler M, Kormann K. Fast matrix-free evaluation of discontinuous Galerkin finite element operators. *ACM Transactions on Mathematical Software* 2019; 45(3): 29:1–29:40. doi: 10.1145/3325864
47. Biehler J, Nitzler J, Wall WA, Gravemeier V. QUEENS – A Software Platform for Uncertainty Quantification, Physics-Informed Machine Learning, Bayesian Optimization, Inverse Problems and Simulation Analytics: User Guide. *AdCo Engineering^{GW}* 2019.
48. Schäfer M, Turek S, Durst F, Krause E, Rannacher R. Benchmark computations of laminar flow around a cylinder. In: *Flow simulation with high-performance computers II*. Springer. 1996 (pp. 547–566)
49. Fehn N, Wall WA, Kronbichler M. On the stability of projection methods for the incompressible Navier–Stokes equations based on high-order discontinuous Galerkin discretizations. *Journal of Computational Physics* 2017; 351: 392–421. doi: <http://dx.doi.org/10.1016/j.jcp.2017.09.031>
50. Fehn N, Wall WA, Kronbichler M. Robust and efficient discontinuous Galerkin methods for under-resolved turbulent incompressible flows. *Journal of Computational Physics* 2018; 372: 667 - 693. doi: <https://doi.org/10.1016/j.jcp.2018.06.037>
51. Ahn S, Fessler JA. Standard errors of mean, variance, and standard deviation estimators. *EECS Department, The University of Michigan* 2003: 1–2.
52. Gerstenberger A, Wall WA. An embedded Dirichlet formulation for 3D continua. *International Journal for Numerical Methods in Engineering* 2010; 82(5): 537–563. doi: <http://dx.doi.org/10.1002/nme.2755>
53. Wall WA, et al. . BACI: A multiphysics simulation environment [Technical Report]. *Munich, Germany: Institute for Computational Mechanics, Technical University of Munich* 2019.
54. Verdugo F, Wall WA. Unified computational framework for the efficient solution of n-field coupled problems with monolithic schemes. *Computer Methods in Applied Mechanics and Engineering* 2016; 310: 335–366. doi: <http://dx.doi.org/10.1016/j.cma.2016.07.016>
55. Klöppel T, Popp A, Küttler U, Wall WA. Fluid–structure interaction for non-conforming interfaces based on a dual mortar formulation. *Computer Methods in Applied Mechanics and Engineering* 2011; 200(45-46): 3111–3126. doi: <http://dx.doi.org/10.1016/j.cma.2011.06.006>
56. Souza Neto dE, Perić D, Dutko M, Owen D. Design of simple low order finite elements for large strain analysis of nearly incompressible solids. *International Journal of Solids and Structures* 1996; 33(20-22): 3277–3296.

How to cite this article: J. Nitzler, J. Biehler, N. Fehn, P.-S. Koutsourelakis, and W. A. Wall (2019), A Generalized Probabilistic Learning Approach for Multi-Fidelity Uncertainty Propagation in Complex Physical Simulations