



HAL
open science

Multilevel Surrogate-based Control Variates

Mohamed Reda El Amri, Paul Mycek, Sophie Ricci, Matthias De Lozzo

► **To cite this version:**

Mohamed Reda El Amri, Paul Mycek, Sophie Ricci, Matthias De Lozzo. Multilevel Surrogate-based Control Variates. IRT Saint-Exupéry; Cerfacs. 2024. <hal-04129561v2>

HAL Id: hal-04129561

<https://hal.science/hal-04129561v2>

Submitted on 18 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Multilevel Surrogate-based Control Variates*

Reda El Amri[†] Paul Mycek^{‡§} Sophie Ricci^{‡§} Matthias De Lozzo[§]

Technical report

Abstract

Estimating statistical parameters of the output of a numerical simulator with random inputs by crude Monte Carlo sampling can be inaccurate when simulations are computationally expensive, so variance reduction techniques must be considered, such as the method of control variates (CVs) or the multilevel Monte Carlo (MLMC) method. In this paper, we propose three multifidelity variance reduction strategies relying on surrogate-based CVs, possibly combined with MLMC. MLCV extends the CV method by using multiple auxiliary random variables devised from surrogate models for simulators of different levels. MLMC-CV improves the MLMC estimator by using a CVs based on a surrogate of the correction term at each level. Further variance reduction is achieved by using the surrogate-based CVs of all the levels in the MLMC-MLCV strategy. Alternative solutions that reduce the subset of surrogates used for the multilevel estimation are also introduced. These strategies are validated on an uncertain 1D heat equation test case from the MLMC literature, where the statistic is the mean integrated temperature at a given time. The results are assessed in terms of the accuracy and computational cost of the multilevel estimators. It is shown that when the lower fidelity outputs are strongly correlated with the high-fidelity outputs, a significant variance reduction is obtained when using surrogate models for the coarser levels only. Taking advantage of pre-existing surrogate models proves to be an even more efficient strategy.

Key words: Multifidelity, multilevel Monte Carlo, control variates, surrogate models, polynomial chaos, variance reduction.

1 Introduction

In recent years, the propagation of uncertainties in numerical simulators has become an essential step in the study of physical phenomena. Therefore, uncertainty quantification (UQ) has emerged as an important element in scientific computing [50, 52]. For complex nonlinear systems, the task of quantifying the effect of uncertainties on the simulator behaviour can pose major challenges as closed-form solutions often do not exist. Sampling-based algorithms are considered the default approach when it comes to UQ for complex nonlinear simulators. Monte Carlo (MC) sampling is arguably the most popular and flexible method. It consists of extracting statistical information from a sample of simulator responses. Due to its non-intrusive nature, MC is straightforward to implement and provides unbiased estimators for classical statistical parameters, e.g., the expectation and the variance. Furthermore, the convergence of MC estimators in terms of their root-mean-square error

*This work was funded by the PIA framework (CGI, ANR) and the industrial members of the IRT Saint Exupéry project R-Evol: Airbus, Liebherr, Altran Technologies, Capgemini DEMS France, CENAERO and Cerfacs.

[†]IFP Énergies Nouvelles, Solaize, France (mohamed-reda.el-amri@ifpen.fr).

[‡]CECI, CNRS - Cerfacs, Toulouse, France (mycek@cerfacs.fr, ricci@cerfacs.fr).

[§]IRT Saint Exupéry, Toulouse, France (matthias.delozzo@irt-saintexupery.com).

(RMSE) as a function of the sample size n is independent of the stochastic input dimension, but, on the other hand, it is very slow, namely at a rate of $n^{-0.5}$. For computationally expensive simulators, the cost of MC is often considered impractically high. In some cases, slight improvements can be obtained through the use of importance sampling [27], latin hypercube sampling [35] or quasi-Monte Carlo (QMC) [13].

Another well-known approach in UQ consists in replacing the simulator by a surrogate model. Fast-to-evaluate surrogates can be used to approximate the high-fidelity model response and therefore reduce drastically the computational cost of estimating statistics. Polynomial chaos (PC) [34, 37], Gaussian process models [45], radial basis functions [41] and (deep) neural networks [54] are commonly used surrogate models. The downside of the surrogate-based approach is that it introduces approximation error, which causes biases in the statistics estimators. Besides, this approximation error tends to increase in high dimension.

One recent MC sampling framework based on control variates (CV) [29, 30, 31, 33] has been extensively developed and used. In a sampling-based CV strategy, one seeks to reduce the variance of the MC estimator of a random variable, arising from the high-fidelity model, by exploiting its correlation with an auxiliary random variable that arises from low-fidelity models approximating the same input-output relationship. In classic CV theory, the mean of the auxiliary random variable is assumed to be known. Unfortunately, in many cases such an assumption is not valid. This creates the need to use another estimator for the auxiliary random variable [42, 39, 43, 44, 24, 48, 49, 47], which involves an additional computational cost. Recently, the adoption of surrogate models as such auxiliary random variables, in particular PC models, has been explored in [16, 56, 17, 25]. The benefits are twofold: (1) the prediction of the surrogate models may be highly correlated with the output high-fidelity model, leading to a reduced variance of the CV estimator as compared to the MC estimation; (2) certain surrogate models provide exact statistics that are needed by the CV approach and for the others, these statistics can be approximated at negligible cost with arbitrary precision.

Another variance reduction technique, called the multilevel MC (MLMC) method [21, 10, 22], uses a hierarchy of models. Originally devised for the estimation of expected values, MLMC has since been extended to the estimation of other statistics, see, e.g., [4, 5] for the estimation of variance and higher order central moments and [36] for the computation of Sobol' indices. Typically, multilevel methods are based on a sequence of levels which correspond to a hierarchy of simulators with increasing accuracy and cost. From a practical standpoint, the different levels often correspond to simulators with increasing mesh resolutions. This translates into a lower accuracy of the so-called *coarse* levels, whereas the *finer* levels correspond to accurate simulators. By construction, MLMC results in an unbiased estimator. It relies on a telescopic sum of terms based on the differences between the successive simulators. It debiases the MC estimator associated with the lowest-level simulator. Many other unbiased multilevel estimators have been devised in recent years. The multi-index MC estimator [26] is an extension of the MLMC estimator such that the telescoping sum idea is used in multiple directions. In [23], the authors developed a variant of MLMC which uses QMC samples instead of independent MC samples for each level.

In this work, we propose new estimation techniques to quantify efficiently the output uncertainties of simulators with limited computational budgets. These approaches build on existing statistical techniques to combine the best of them: MC sampling to guarantee the unbiasedness of the estimators and the independence of their accuracy with respect to the input dimension, multilevel estimation to reduce their variance by exploiting some hierarchical fidelity structure, and surrogate-based CVs to further reduce the variance, especially on the coarser levels, as efficiently as possible. It is worth noting that we do not seek to estimate statistics from accurate surrogate models, but rather from sampling techniques assisted by multifidelity simulators and coarse-to-fine

surrogate models. A first strategy, named *multilevel control variates* (MLCV), combines multiple CVs based on surrogate models of simulators with different fidelity levels. Although this strategy does not build on the MLMC approach specifically, it still exploits multilevel information through the surrogates constructed at different levels. The second strategy, named *multilevel Monte Carlo with control variates* (MLMC-MLCV), utilizes CVs in an information fusion framework to exploit synergies between the flexible MLMC sampling and the correlation shared between the high- and low-fidelity components. The numerical results show that the unbiased MLMC-MLCV estimators can converge faster than the existing estimators.

This technical report is organized as follows. We introduce notations and the necessary mathematical background in section 2. In section 3, we briefly discuss the MLMC estimator and then define our proposed MLMC-MLCV estimator that combines multilevel sampling and surrogate-based CV. In section 4, we conduct numerical experiments to support the theoretical results. Section 5 proposes concluding remarks.

2 Background

In this section, we summarize the important results of statistical estimation using control variates and we show how surrogate models can be leveraged in this setting.

2.1 Notation

We first introduce a few notations. Throughout the rest of this paper, the high-fidelity numerical model is abstractly represented by the deterministic mapping

$$\begin{aligned} f : \Xi &\longrightarrow \mathbb{R} \\ \mathbf{x} &\longmapsto f(\mathbf{x}), \end{aligned} \tag{2.1}$$

where $\mathbf{x} := (x_1 \ \cdots \ x_d)^\top$ is a vector of d uncertain input parameters evolving in a measurable space denoted by $\Xi = \Xi_1 \times \cdots \times \Xi_d$, with $\Xi_i \subset \mathbb{R}$. Following a probabilistic approach, the d -dimensional input is modelled as a random vector $\mathbf{X} : \Omega \rightarrow \mathbb{R}^d$ with known *joint* probability distribution. In this work, we seek an accurate estimator of some statistic θ of the output random variable $Y := f(\mathbf{X}) : \Omega \rightarrow \mathbb{R}$, e.g., its expected value ($\theta = \mathbb{E}[Y]$) or variance ($\theta = \mathbb{V}[Y]$), at a reasonable computational cost \mathcal{C} .

2.2 Crude Monte Carlo

In practice, the MC estimator $\hat{\theta}$ of the statistic θ is based on n observations $Y^{(1)}, \dots, Y^{(n)}$ defined as $Y^{(i)} = f(\mathbf{X}^{(i)})$ where $\mathcal{X} = \{\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(n)}\}$ is a n -sample of \mathbf{X} , that is, a collection of independent and identically distributed random variables with the same distribution as \mathbf{X} . For instance, the *sample mean* and (unbiased) *sample variance* estimators,

$$\hat{E}[Y] = n^{-1} \sum_{i=1}^n Y^{(i)} \quad \text{and} \quad \hat{V}[Y] = (n-1)^{-1} \sum_{i=1}^n (Y^{(i)} - \hat{E}[Y])^2 \tag{2.2}$$

are MC estimators of the expectation and variance of Y , respectively. It is well known that these estimators are *unbiased*, that is, $\mathbb{E}[\hat{\theta}] = \theta$, so that the mean square error (MSE) of $\hat{\theta}$ reduces to the variance of the estimator:

$$\text{MSE}(\hat{\theta}, \theta) := \mathbb{E}[(\hat{\theta} - \theta)^2] = \mathbb{V}[\hat{\theta}] + (\mathbb{E}[\hat{\theta}] - \theta)^2 = \mathbb{V}[\hat{\theta}]. \tag{2.3}$$

The convergence of these MC estimators is known to be slow, so that variance reduction techniques are needed when dealing with computationally expensive simulators.

2.3 Control variates

In this section, we present a well-known variance reduction technique using auxiliary random variables Z_1, \dots, Z_M as control variates. We denote by τ_1, \dots, τ_M statistical parameters related to the control variates Z_1, \dots, Z_M , and assume they are known exactly. Then, the CV estimator is defined as

$$\hat{\theta}^{\text{CV}}(\boldsymbol{\alpha}) = \hat{\theta} - \boldsymbol{\alpha}^\top(\hat{\boldsymbol{\tau}} - \boldsymbol{\tau}), \quad (2.4)$$

where $\hat{\theta}$ and $\hat{\boldsymbol{\tau}} = (\hat{\tau}_1 \ \dots \ \hat{\tau}_M)^\top$ are *unbiased* MC estimators of θ and $\boldsymbol{\tau} = (\tau_1 \ \dots \ \tau_M)^\top$, respectively, based on a *common* input n -sample \mathcal{X} , and where $\boldsymbol{\alpha} \in \mathbb{R}^M$ is the control parameter. We note that the CV estimator is unbiased by construction, regardless of the value of the parameter $\boldsymbol{\alpha}$, and that its variance reads

$$\mathbb{V}[\hat{\theta}^{\text{CV}}(\boldsymbol{\alpha})] = \mathbb{V}[\hat{\theta}] + \boldsymbol{\alpha}^\top \boldsymbol{\Sigma} \boldsymbol{\alpha} - 2\boldsymbol{\alpha}^\top \mathbf{c}. \quad (2.5)$$

with $\mathbf{c} := \mathbb{C}[\hat{\boldsymbol{\tau}}, \hat{\theta}] \in \mathbb{R}^M$ and $\boldsymbol{\Sigma} := \mathbb{C}[\hat{\boldsymbol{\tau}}] \in \mathbb{R}^{M \times M}$. To fully take advantage of the control variates, the parameter $\boldsymbol{\alpha}$ is selected so as to minimize the variance (2.5) of the CV estimator. Assuming that the covariance matrix $\boldsymbol{\Sigma}$ is non-singular (and thus symmetric positive definite, SPD), the first- and second-order optimality conditions of the minimization problem imply that there exists a unique optimal solution,

$$\boldsymbol{\alpha}^* = \boldsymbol{\Sigma}^{-1} \mathbf{c}. \quad (2.6)$$

The optimal CV estimator is thus $\hat{\theta}^{\text{CV}}(\boldsymbol{\alpha}^*)$, and its variance is given by

$$\mathbb{V}[\hat{\theta}^{\text{CV}}(\boldsymbol{\alpha}^*)] = (1 - R^2) \mathbb{V}[\hat{\theta}], \quad (2.7)$$

where $R^2 = \mathbb{V}[\hat{\theta}]^{-1} \mathbf{c}^\top \boldsymbol{\Sigma}^{-1} \mathbf{c}$ is nonnegative, as $\boldsymbol{\Sigma}^{-1}$ is SPD. Denoting by $\mathbf{D} = \text{Diag}(\boldsymbol{\Sigma}) \in \mathbb{R}^{M \times M}$ the diagonal matrix consisting of the diagonal of $\boldsymbol{\Sigma}$, it can be shown further that $R^2 = \mathbf{r}^\top \mathbf{R}^{-1} \mathbf{r}$, where $\mathbf{r} = (\mathbb{V}[\hat{\theta}] \mathbf{D})^{-1/2} \mathbf{c}$ is the vector of the Pearson correlation coefficients between $\hat{\theta}$ and $\hat{\boldsymbol{\tau}}$, and $\mathbf{R} = \mathbf{D}^{-1/2} \boldsymbol{\Sigma} \mathbf{D}^{-1/2}$ is the correlation matrix of $\hat{\boldsymbol{\tau}}$. Thus, $R^2 \in [0, 1]$ corresponds to the squared coefficient of multiple correlation between $\hat{\theta}$ and the elements of $\hat{\boldsymbol{\tau}}$ [1, section 2.5.2]. Consequently, $0 \leq \mathbb{V}[\hat{\theta}^{\text{CV}}(\boldsymbol{\alpha}^*)] \leq \mathbb{V}[\hat{\theta}]$, and we refer to R^2 as the variance reduction factor of the CV estimator. This shows that the variance of the optimal CV estimator $\hat{\theta}^{\text{CV}}(\boldsymbol{\alpha}^*)$ is always reduced (or, rigorously speaking, not increased) compared to the MC estimator $\hat{\theta}$. Furthermore, the higher R^2 , the greater the reduction in variance.

Remark 2.1. *The requirement that $\boldsymbol{\Sigma}$ be non-singular is a reasonable one. Indeed, let us suppose that $\boldsymbol{\Sigma}$ is singular. Then, because $\boldsymbol{\Sigma}$ is positive semi-definite by construction, this implies that there exists a nonzero vector $\boldsymbol{\eta} \in \mathbb{R}^M \setminus \{\mathbf{0}\}$ such that $\boldsymbol{\eta}^\top \boldsymbol{\Sigma} \boldsymbol{\eta} = \mathbb{V}[\boldsymbol{\eta}^\top \hat{\boldsymbol{\tau}}] = 0$, indicating that any one element of $\hat{\boldsymbol{\tau}}$ can be expressed as an affine function of the others. As such, it does not bring any additional information to the CV estimator, so that at least one of the control variates can simply be discarded.*

A desirable property of the CV estimator is that increasing the number of control variates improves the CV estimator. Specifically, under mild assumptions, proposition 2.1 states that the variance of the CV estimator is reduced (or, rigorously speaking, not increased) when adding a new control variate.

Proposition 2.1. *Let $\hat{\boldsymbol{\tau}}_+ := [\hat{\boldsymbol{\tau}}^\top \ \hat{\tau}_{M+1}]^\top$, $\boldsymbol{\tau}_+ := [\boldsymbol{\tau}^\top \ \tau_{M+1}]^\top$, and define $\mathbf{c}_+ = \mathbb{C}[\hat{\boldsymbol{\tau}}_+, \hat{\theta}]$, and $\boldsymbol{\Sigma}_+ = \mathbb{C}[\hat{\boldsymbol{\tau}}_+]$. We further assume that $\boldsymbol{\Sigma}$ and $\boldsymbol{\Sigma}_+$ are non-singular and that $\mathbb{V}[\hat{\tau}_{M+1}] > 0$. Let $\hat{\theta}^{\text{CV}}(\boldsymbol{\alpha}_+^*) := \hat{\theta} - \boldsymbol{\alpha}_+^{*\top}(\hat{\boldsymbol{\tau}}_+ - \boldsymbol{\tau}_+)$, with $\boldsymbol{\alpha}_+^* = \boldsymbol{\Sigma}_+^{-1} \mathbf{c}_+$, be the optimal CV estimator based on $M + 1$ control variates, and let R_+^2 denote its variance reduction factor. Then $R_+^2 \geq R^2$.*

Proof. We have $R_+^2 = \mathbf{r}_+^\top \mathbf{R}_+^{-1} \mathbf{r}_+$, where $\mathbf{r}_+ = [\mathbf{r}^\top \ \gamma]^\top$ and $\mathbf{R}_+ = \begin{bmatrix} \mathbf{R} & \mathbf{u} \\ \mathbf{u}^\top & 1 \end{bmatrix}$, and where $\mathbf{u} = (\mathbb{V}[\hat{\tau}_{M+1}] \mathbf{D})^{-1/2} \mathbb{C}[\hat{\tau}, \hat{\tau}_{M+1}]$ and $\gamma = (\mathbb{V}[\hat{\theta}] \mathbb{V}[\hat{\tau}_{M+1}])^{-1/2} \mathbb{C}[\hat{\theta}, \hat{\tau}_{M+1}]$. Note that \mathbf{R} and \mathbf{R}_+ are both non-singular, because so are $\mathbf{\Sigma}$ and $\mathbf{\Sigma}_+$. The augmented matrix \mathbf{R}_+ may be inverted by block using the Schur complement $s = 1 - \mathbf{u}^\top \mathbf{R}^{-1} \mathbf{u} \neq 0$ of the (1,1)-block \mathbf{R} . It follows that $R_+^2 = R^2 + s^{-1}(\gamma - \mathbf{u}^\top \mathbf{R}^{-1} \mathbf{r})^2$. Now, because \mathbf{R}_+ is SPD (it is positive semidefinite by construction and non-singular by assumption), $\mathbf{x}^\top \mathbf{R}_+ \mathbf{x} > 0$ for any choice of $\mathbf{x} \neq \mathbf{0}$. The particular choice of $\mathbf{x} = [-\mathbf{u}^\top \mathbf{R}^{-1} \ 1]^\top$ implies that $s > 0$, which in turn implies that $R_+^2 \geq R^2$. \square

Corollary 2.2. *Under the assumptions of proposition 2.1, the equality $R_+^2 = R^2$ holds if and only if $\gamma = \mathbf{u}^\top \mathbf{R}^{-1} \mathbf{r}$.*

Proof. Straightforward from $R_+^2 = R^2 + s^{-1}(\gamma - \mathbf{u}^\top \mathbf{R}^{-1} \mathbf{r})^2$ with $s > 0$. \square

For the CV estimation of specific statistics, the expressions of \mathbf{c} and $\mathbf{\Sigma}$ may be further reduced to involve statistics on Y and $\mathbf{Z} := (Z_1 \ \cdots \ Z_M)^\top$ directly. These expressions, as well as the consequences on eqs. (2.6) and (2.7), are given in appendix A for the CV estimators of $\mathbb{E}[Y]$ and $\mathbb{V}[Y]$.

It should be noted at this point that, in practice, the optimal parameter $\boldsymbol{\alpha}^*$ needs to be approximated. Specifically, the statistics involved in \mathbf{c} and $\mathbf{\Sigma}$ need to be estimated. This can be done either using the same sample \mathcal{X} as for the CV estimator itself, or using an independent, pilot sample \mathcal{X}' . The former strategy introduces a bias in the resulting CV estimator, while the latter guarantees unbiasedness but requires additional high-fidelity simulator evaluations, so that the former strategy is generally preferred. Furthermore, statistical remedies such as jackknifing, splitting or bootstrapping have been proposed to reduce or eliminate the bias introduced by the former strategy [38]. In both strategies, however, neither the theoretical variance reduction given by eq. (2.7) nor proposition 2.1 hold anymore.

Remark 2.2. *In some instances, because it only involves the control variates, $\mathbf{\Sigma}$ may be known exactly (see remark 3.1 for an example in a multilevel setting involving PC-based control variates) or estimated accurately using many independent samples at negligible cost. Unfortunately, this is not the case for \mathbf{c} , which involves the output $Y = f(\mathbf{X})$ of the high-fidelity simulator f .*

2.4 Surrogate-based control variates

The efficiency of the CV approach relies on the strong correlation between $Y = f(\mathbf{X})$ and the control variates \mathbf{Z} . In a multifidelity framework, the control variates correspond to the output of low-fidelity versions of f , typically in the form of simulators with simplified physics and/or coarser discretization. In many applications, the exact statistical measures $\boldsymbol{\tau}$ of such control variates \mathbf{Z} may not be available. One way to circumvent this limitation is to use additional samples to also estimate $\boldsymbol{\tau}$. This type of estimators led to an *approximate* CV class of methods [24, 39, 43], where different model management and sample allocation strategies may be used to find a suitable trade-off between the additional cost of sampling \mathbf{Z} and the resulting variance reduction. Recently, an optimal strategy was proposed [48, 49, 47] leading to the so-called multilevel best linear unbiased estimator (MLBLUE), as well as a unifying framework for a large class of multilevel and multifidelity estimators. However, ACV estimators typically yield lower variance reduction than their exact CV counterparts, as illustrated in appendix B on bi-fidelity (A)CV estimators of the expectation.

In this paper, we focus on the case where the low-fidelity models correspond to surrogate models of the high-fidelity simulator f , which allows us to use the original CV approach described in

section 2.3. Indeed, compared to the cost of evaluating the high-fidelity simulator f , the cost of evaluating surrogate models can be neglected, so that the statistical parameters $\boldsymbol{\tau}$ of the output of the surrogate models can be estimated arbitrarily accurately at negligible cost. In fact, in some instances, the exact statistical parameters $\boldsymbol{\tau}$ may even be directly available, as detailed hereafter. As a direct consequence, as discussed above, the question of sample allocation need not be addressed, and greater variance reduction is expected. Using surrogate models in a CV strategy has been explored previously in [16, 56], where the available computational budget is allocated both to the construction of the surrogates and to the actual CV estimations. In [56], the authors introduce an approach that optimally balances the computational effort needed to select the optimal degree of the polynomial chaos (PC) expansion used in a stochastic Galerkin CV approach. In our work, we focus on the situation in which the surrogate models are available, so that we do not consider any optimization strategy for the construction of the surrogates. Nevertheless, for the fairness of comparison, we still report the pre-processing cost of constructing the surrogate models.

We now briefly describe three different surrogate models commonly used in a UQ framework, and discuss the availability of statistics of their output.

2.4.1 Gaussian process modelling

We assume that the simulator f is a realization of a Gaussian process (GP) F indexed by \mathbf{x} and defined by its mean function m_F and covariance kernel k_F ,

$$\mathbb{E}[F(\mathbf{x})] = m_F(\mathbf{x}), \quad (2.8)$$

$$\mathbb{C}[F(\mathbf{x}), F(\mathbf{x}')] = k_F(\mathbf{x}, \mathbf{x}'), \quad \forall \mathbf{x}, \mathbf{x}' \in \Xi. \quad (2.9)$$

In practice, one can parametrize the forms of m_F and k_F . For example, in the widely used ordinary GP method, a stationary GP is assumed. In this case, m_F is set as a constant $m_F(\mathbf{x}) = m$. More importantly, it is assumed that $k_F(\mathbf{x}, \mathbf{x}') = \bar{k}_F(\mathbf{x} - \mathbf{x}')$, and $k_F(\mathbf{x}, \mathbf{x}) = \bar{k}_F(0) = \sigma^2$ is a constant. Popular forms of kernels include polynomial, exponential, Gaussian, and Matérn functions. For example, the Gaussian kernel can be written as $k_F(\mathbf{x}, \mathbf{x}') = \sigma^2 \exp(-\frac{1}{2}\|\mathbf{x} - \mathbf{x}'\|_{\mathbf{h}}^2)$, where the weighted norm is defined as $\|\mathbf{x} - \mathbf{x}'\|_{\mathbf{h}} = \left(\sum_{i=1}^d \frac{(x_i - x'_i)^2}{h_i^2}\right)^{1/2}$ where h_1, \dots, h_d are correlation lengths. The hyperparameters σ and h_i can be obtained by maximum likelihood. Then, given n observations $\mathcal{F} = (f(\mathbf{x}^{(1)}) \ \dots \ f(\mathbf{x}^{(n)}))^{\top}$ of F at $\mathcal{X} = (\mathbf{x}^{(1)} \ \dots \ \mathbf{x}^{(n)})^{\top}$, the posterior \tilde{F} of F can be defined as

$$\tilde{F} = [F \mid F(\mathcal{X}) = \mathcal{F}], \quad (2.10)$$

whose expectation and covariance are given by

$$m_{\tilde{F}}(\mathbf{x}) = \mathbb{E}[\tilde{F}(\mathbf{x})] = m_F(\mathbf{x}) + k_F(\mathbf{x}, \mathcal{X})^{\top} k_F(\mathcal{X}, \mathcal{X})^{-1} (\mathcal{F} - m_F(\mathcal{X})), \quad (2.11)$$

$$k_{\tilde{F}}(\mathbf{x}, \mathbf{x}') = \mathbb{C}[\tilde{F}(\mathbf{x}), \tilde{F}(\mathbf{x}')] = k_F(\mathbf{x}, \mathbf{x}') - k_F(\mathbf{x}, \mathcal{X})^{\top} k_F(\mathcal{X}, \mathcal{X})^{-1} k_F(\mathcal{X}, \mathbf{x}'). \quad (2.12)$$

Thereafter, the high-fidelity model f at \mathbf{x} will be approximated by the conditional expectation,

$$g^{\text{GP}}(\mathbf{x}) = m_{\tilde{F}}(\mathbf{x}). \quad (2.13)$$

Thus, the expectation and variance of $g^{\text{GP}}(\mathbf{X})$ are defined as

$$\mathbb{E}[g^{\text{GP}}(\mathbf{X})] = \int_{\Xi} m_{\tilde{F}}(\mathbf{x}) p_{\mathbf{X}}(\mathbf{x}) \, d\mathbf{x}, \quad (2.14)$$

$$\mathbb{V}[g^{\text{GP}}(\mathbf{X})] = \int_{\Xi} (m_{\tilde{F}}(\mathbf{x}) - \mathbb{E}[g^{\text{GP}}(\mathbf{X})])^2 p_{\mathbf{X}}(\mathbf{x}) \, d\mathbf{x}. \quad (2.15)$$

These two statistics can be approximated empirically by taking a large sample of \mathbf{X} , as the GP model is inexpensive. Analytical formulas exist for some pairs of distributions and covariance functions (see [8, Table 1] and [28, 15]).

2.4.2 Taylor polynomials

We assume that the numerical simulator f is infinitely differentiable and that the moments of \mathbf{X} are finite. Under this assumption, it is possible to expand the original simulator f around the input's expected value $\boldsymbol{\mu}_{\mathbf{X}} := \mathbb{E}[\mathbf{X}]$ as the infinite polynomial series according to Taylor's theorem,

$$f(\mathbf{X}) = g^{\text{T}\infty}(\mathbf{X}) = \sum_{|\boldsymbol{\beta}| \leq p} \frac{(\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}})^{\boldsymbol{\beta}}}{\boldsymbol{\beta}!} D^{\boldsymbol{\beta}} f(\boldsymbol{\mu}_{\mathbf{X}}), \quad (2.16)$$

where $\boldsymbol{\beta} \in \mathbb{N}^d$, $|\boldsymbol{\beta}| = \sum_{i=1}^d \beta_i$, $\boldsymbol{\beta}! = \prod_{i=1}^d \beta_i!$, $x^{\boldsymbol{\beta}} = \prod_{i=1}^d x_i^{\beta_i}$ and $D^{\boldsymbol{\beta}} f = \frac{\partial^{|\boldsymbol{\beta}|} f}{\partial \beta_1 x_1 \dots \partial \beta_d x_d}$. Thus, the function f may be approximated by the first- and second-order Taylor polynomials,

$$g^{\text{T}1}(\mathbf{X}) = f(\boldsymbol{\mu}_{\mathbf{X}}) + \mathbf{J}_f(\boldsymbol{\mu}_{\mathbf{X}})(\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}}), \quad (2.17)$$

$$g^{\text{T}2}(\mathbf{X}) = f(\boldsymbol{\mu}_{\mathbf{X}}) + \mathbf{J}_f(\boldsymbol{\mu}_{\mathbf{X}})(\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}}) + \frac{1}{2}(\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}})^{\text{T}} \mathbf{H}_f(\boldsymbol{\mu}_{\mathbf{X}})(\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}}), \quad (2.18)$$

where $\mathbf{J}_f(\boldsymbol{\mu}_{\mathbf{X}}) = \nabla f(\boldsymbol{\mu}_{\mathbf{X}})^{\text{T}} \in \mathbb{R}^{1 \times d}$ and $\mathbf{H}_f(\boldsymbol{\mu}_{\mathbf{X}}) \in \mathbb{R}^{d \times d}$ are the Jacobian and Hessian matrices of f at $\boldsymbol{\mu}_{\mathbf{X}}$, respectively. In practical computations, the derivatives may be approximated by numerical differentiation if they are not provided. The expectation and variance of $g^{\text{T}1}(\mathbf{X})$ and $g^{\text{T}2}(\mathbf{X})$ are then defined as

$$\mathbb{E}[g^{\text{T}1}(\mathbf{X})] = f(\boldsymbol{\mu}_{\mathbf{X}}), \quad (2.19)$$

$$\mathbb{V}[g^{\text{T}1}(\mathbf{X})] = \mathbf{J}_f(\boldsymbol{\mu}_{\mathbf{X}})^{\odot 2} \boldsymbol{\sigma}_{\mathbf{X}}^2, \quad (2.20)$$

$$\mathbb{E}[g^{\text{T}2}(\mathbf{X})] = f(\boldsymbol{\mu}_{\mathbf{X}}) + \frac{1}{2} \text{Tr}(\mathbf{H}_f(\boldsymbol{\mu}_{\mathbf{X}}) \boldsymbol{\Sigma}_{\mathbf{X}}^2), \quad (2.21)$$

$$\mathbb{V}[g^{\text{T}2}(\mathbf{X})] = \text{Tr}(\mathbf{J}_f(\boldsymbol{\mu}_{\mathbf{X}})^{\text{T}} \mathbf{J}_f(\boldsymbol{\mu}_{\mathbf{X}}) \boldsymbol{\Sigma}_{\mathbf{X}}^2) + \frac{1}{2} \text{Tr}(\mathbf{H}_f(\boldsymbol{\mu}_{\mathbf{X}})^{\odot 2} \boldsymbol{\sigma}_{\mathbf{X}}^2 (\boldsymbol{\sigma}_{\mathbf{X}}^2)^{\text{T}}), \quad (2.22)$$

where for any matrix \mathbf{A} , $\mathbf{A}^{\odot 2}$ denotes the element-wise square of \mathbf{A} , $\text{Tr}(\mathbf{A})$ is the trace of \mathbf{A} , $\boldsymbol{\sigma}_{\mathbf{X}}^2 = [\sigma_{X_1}^2 \ \dots \ \sigma_{X_d}^2]^{\text{T}} \in \mathbb{R}^d$ is the element-wise variance of \mathbf{X} , and $\boldsymbol{\Sigma}_{\mathbf{X}}^2 = \text{Diag}(\boldsymbol{\sigma}_{\mathbf{X}}^2) \in \mathbb{R}^{d \times d}$.

2.4.3 Polynomial chaos expansion

We consider the truncated polynomial chaos (PC) expansion of $f(\mathbf{X})$ [55, 9, 20, 32],

$$f(\mathbf{X}) \simeq g^{\text{PC}P}(\mathbf{X}) = \sum_{k=0}^P \mathbf{g}_k \Psi_k(\mathbf{X}), \quad (2.23)$$

where the coefficients \mathbf{g}_k are real scalars and Ψ_k are orthonormal multivariate polynomials:

$$\forall i, j \geq 0, \quad \langle \Psi_i, \Psi_j \rangle_{p_{\mathbf{X}}} := \int_{\Xi} \Psi_i(\mathbf{x}) \Psi_j(\mathbf{x}) p_{\mathbf{X}}(\mathbf{x}) \, d\mathbf{x} = \mathbb{E}[\Psi_i(\mathbf{X}) \Psi_j(\mathbf{X})] = \delta_{ij}, \quad (2.24)$$

where δ_{ij} denotes the Kronecker delta. The coefficients \mathbf{g}_k may be approximated using non-intrusive techniques such as non-intrusive (pseudo)spectral projection [46, 12, 11], regression [3, 51, 6, 7],

interpolation [2, 40], or from intrusive approaches such as the stochastic Galerkin method [20, 32]. Assuming $\Psi_0 \equiv 1$, the expectation and variance of $g^{\text{PCP}}(\mathbf{X})$ are then given by

$$\mathbb{E}[g^{\text{PCP}}(\mathbf{X})] = \mathbf{g}_0 \quad \text{and} \quad \mathbb{V}[g^{\text{PCP}}(\mathbf{X})] = \sum_{k=1}^P \mathbf{g}_k^2. \quad (2.25)$$

3 Multilevel estimators

In this section, we present so-called *multilevel* statistical estimation techniques based on a sequence of simulators $(f_\ell)_{\ell=0}^L$, with increasing accuracy and computational cost, where $f_L \equiv f$ denotes the high-fidelity numerical simulator. The levels are ordered from the *coarsest* ($\ell = 0$) to the *finest* ($\ell = L$). We denote by Y_ℓ the random variable $Y_\ell = f_\ell(\mathbf{X})$ and θ_ℓ the statistic of $f_\ell(\mathbf{X})$ increasingly close to $\theta_L \equiv \theta$.

3.1 Multilevel Monte Carlo

The statistic θ_L can be expressed as the telescoping sum

$$\theta_L = \sum_{\ell=0}^L T_\ell, \quad (3.1)$$

where $T_\ell = \theta_\ell - \theta_{\ell-1}$, and, by convention, $\theta_{-1} := 0$. The MLMC estimator $\hat{\theta}_L^{\text{MLMC}}$ of θ_L is then defined as [21, 22, 53]

$$\hat{\theta}^{\text{MLMC}} = \sum_{\ell=0}^L \hat{T}_\ell^{(\ell)}, \quad (3.2)$$

where, at each level ℓ , $\hat{T}_\ell^{(\ell)}$ is an unbiased MC estimator of T_ℓ , based on an input sample $\mathcal{X}^{(\ell)} = \{\mathbf{X}^{(\ell,i)}\}_{i=1}^{n_\ell}$ such that the members of $\mathcal{X}^{(\ell)}$ and $\mathcal{X}^{(\ell')}$ are mutually independent for $\ell \neq \ell'$. In many instances, $\hat{T}_\ell^{(\ell)}$ is actually defined as $\hat{T}_\ell^{(\ell)} = \hat{\theta}_\ell^{(\ell)} - \hat{\theta}_{\ell-1}^{(\ell)}$, where $\hat{\theta}_k^{(\ell)}$ denotes the unbiased MC estimator of θ_k based on the simulator f_k using the n_ℓ -sample $\mathcal{X}^{(\ell)}$.

For example, the MLMC estimator $\hat{E}^{\text{MLMC}}[Y]$ of the expectation $\mathbb{E}[Y]$ is given by

$$\hat{E}^{\text{MLMC}}[Y] = \hat{E}^{(0)}[Y_0] + \sum_{\ell=1}^L \hat{E}^{(\ell)}[Y_\ell] - \hat{E}^{(\ell)}[Y_{\ell-1}], \quad (3.3)$$

where $\hat{E}^{(\ell)}[Y_k] = n_\ell^{-1} \sum_{i=1}^{n_\ell} f_k(\mathbf{X}^{(\ell,i)})$, with $k \in \{\ell, \ell-1\}$. We stress that the correction terms at each level ℓ are computed from the same input sample $\mathcal{X}^{(\ell)}$, but using two successive simulators, f_ℓ and $f_{\ell-1}$. Similarly, the MLMC estimator $\hat{V}^{\text{MLMC}}[Y]$ of the variance $\mathbb{V}[Y]$ is defined as [4]

$$\hat{V}^{\text{MLMC}}[Y] = \hat{V}^{(0)}[Y_0] + \sum_{\ell=1}^L \hat{V}^{(\ell)}[Y_\ell] - \hat{V}^{(\ell)}[Y_{\ell-1}], \quad (3.4)$$

where $\hat{V}^{(\ell)}[Y_k] = \frac{n_\ell}{n_\ell-1} (\hat{E}^{(\ell)}[Y_k^2] - \hat{E}^{(\ell)}[Y_k]^2)$ is the single-level unbiased MC variance estimator. Owing to the independence of the estimators $\hat{T}_\ell^{(\ell)}$, the variance of the MLMC estimator is

$$\mathbb{V}[\hat{\theta}^{\text{MLMC}}] = \sum_{\ell=0}^L \mathbb{V}[\hat{T}_\ell^{(\ell)}]. \quad (3.5)$$

In practice, the MLMC method relies on the allocation of the total (expected) computational cost of the MLMC estimator across the different levels, with

$$\text{cost}(\hat{\theta}^{\text{MLMC}}) = \sum_{\ell=0}^L n_{\ell}(\mathcal{C}_{\ell} + \mathcal{C}_{\ell-1}), \quad (3.6)$$

where \mathcal{C}_{ℓ} is the (expected) cost of one evaluation of the simulator f_{ℓ} , with $\mathcal{C}_{-1} := 0$ by convention. Thus, a key aspect is played by the choice of the number of samples n_{ℓ} allocated to each level ℓ . The goal is to find the sample sizes n_0, \dots, n_L that minimize the variance of the estimator $\mathbb{V}[\hat{\theta}^{\text{MLMC}}]$ given a computational budget \mathcal{C} . Thus, the sample allocation problem reads

$$\begin{aligned} & \underset{n_0, \dots, n_L \in \mathbb{N}^*}{\text{minimize}} && \mathbb{V}[\hat{\theta}^{\text{MLMC}}] \\ & \text{subject to} && \text{cost}(\hat{\theta}^{\text{MLMC}}) = \mathcal{C}. \end{aligned} \quad (3.7)$$

In practice, $\mathbb{V}[\hat{\theta}^{\text{MLMC}}]$ is not known, and we instead rely on the assumption that $\mathbb{V}[\hat{T}_{\ell}^{(\ell)}] \lesssim n_{\ell}^{-1} \mathcal{V}_{\ell}$, with \mathcal{V}_{ℓ} independent of n_{ℓ} [53, 36]. This is a reasonable assumption that holds for the MLMC estimation of the expectation, variance and covariance [36, Table 1]. Note that, for the estimation of the expectation, we have $\mathbb{V}[\hat{T}_{\ell}^{(\ell)}] = n_{\ell}^{-1} \mathcal{V}_{\ell}$, with $\mathcal{V}_{\ell} = \mathbb{V}[Y_{\ell} - Y_{\ell-1}]$ and $Y_{-1} \equiv 0$. The sample allocation problem eq. (3.7) is then replaced with

$$\begin{aligned} & \underset{n_0, \dots, n_L \in \mathbb{N}^*}{\text{minimize}} && \sum_{\ell=0}^L n_{\ell}^{-1} \mathcal{V}_{\ell} \\ & \text{subject to} && \text{cost}(\hat{\theta}^{\text{MLMC}}) = \mathcal{C}, \end{aligned} \quad (3.8)$$

which is equivalent for the expectation, and an approximation for other statistics. This minimization problem has a unique solution which can be computed analytically (see, e.g., [21, 10, 36]).

3.2 Multilevel surrogate-based control variate strategies

In this section, we introduce various surrogate-based control variate strategies in a multilevel framework where a hierarchy of simulators $(f_{\ell})_{\ell=0}^L$ is available. These strategies rely on using the random variables $(g_{\ell}(\mathbf{X}))_{\ell=0}^L$ as control variates, where g_{ℓ} is a surrogate model of the simulator f_{ℓ} .

3.2.1 Multilevel control variates (MLCV)

The first strategy, referred to as *multilevel control variates* and hereafter abbreviated MLCV, consists of using the surrogate models at *all levels* to build the control variates in eq. (2.4). Thus, τ_{ℓ} corresponds to the statistical measure of the random variable $Z_{\ell} = g_{\ell}(\mathbf{X})$, and $\hat{\tau}_{\ell}$ to its unbiased MC estimator. For instance, the MLCV estimator of the expectation based on an n_L -sample $\mathcal{X}^{(L)} = \{\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(n_L)}\}$ reads

$$\hat{E}^{\text{MLCV}}[Y](\boldsymbol{\alpha}) = \hat{E}^{(L)}[Y_L] - \sum_{\ell=0}^L \alpha_{\ell} \left(\hat{E}^{(L)}[Z_{\ell}] - \mu_{Z_{\ell}} \right), \quad (3.9)$$

where $\hat{E}^{(L)}[Y_L] = n_L^{-1} \sum_{i=1}^{n_L} f_L(\mathbf{X}^{(i)})$, $\hat{E}^{(L)}[Z_{\ell}] = n_L^{-1} \sum_{i=1}^{n_L} g_{\ell}(\mathbf{X}^{(i)})$, and $\mu_{Z_{\ell}} = \mathbb{E}[g_{\ell}(\mathbf{X})]$. Note that this approach does not build on the MLMC methodology described in section 3.1, but still exploits multilevel information through the surrogates constructed at different levels.

3.2.2 Multilevel Monte Carlo with control variates (MLMC-CV)

The second strategy consists of improving the MLMC estimator eq. (3.2) using the surrogate-based control variates Z_0, \dots, Z_L . Specifically, the MLMC-CV estimator improves the MC estimation of each of the correction terms of the MLMC estimator by using a surrogate model of the corresponding correction term as a control variate. Note that the MLMF approach proposed in [18, 19] is based on a similar strategy, using arbitrary low-fidelity models at level in an *approximate* CV setting. The MLMC-CV estimator reads

$$\hat{\theta}^{\text{MLMC-CV}}(\alpha_1, \dots, \alpha_L) = \sum_{\ell=0}^L \hat{T}_\ell^{\text{CV}}(\alpha_\ell), \quad (3.10)$$

where $\hat{T}_\ell^{\text{CV}}(\alpha_\ell)$ is the CV estimator of the multilevel correction term T_ℓ (cf. eq. (3.1)),

$$\hat{T}_\ell^{\text{CV}}(\alpha_\ell) = \hat{T}_\ell^{(\ell)} - \alpha_\ell(\hat{U}_\ell^{(\ell)} - U_\ell), \quad (3.11)$$

with $\hat{U}_k^{(\ell)}$ an unbiased MC estimator of the control variate statistic $U_k = \tau_k - \tau_{k-1}$, based on the same input sample $\mathcal{X}^{(\ell)}$ as $\hat{T}_\ell^{(\ell)}$, again with members of $\mathcal{X}^{(\ell)}$ and $\mathcal{X}^{(\ell')}$ being independent for $\ell \neq \ell'$. Note that, in eq. (3.11), because a single control variate is used at each level, the definition of $\hat{U}_k^{(\ell)}$ is used in the specific case where $k = \ell$. The more general definition when k and ℓ are not necessarily equal will be useful later in section 3.2.3, where we consider multiple control variates per level. In practice, $\hat{U}_k^{(\ell)}$ may be defined as $\hat{U}_k^{(\ell)} = \hat{\tau}_k^{(\ell)} - \hat{\tau}_{k-1}^{(\ell)}$, where $\hat{\tau}_k^{(\ell)}$ is an unbiased estimator of τ_k using the n_ℓ -sample $\mathcal{X}^{(\ell)}$.

The optimal value α_ℓ^* for α_ℓ is obtained individually for each $\ell = 0, \dots, L$ as the optimal (single) CV parameter for $\hat{T}_\ell^{\text{CV}}(\alpha_\ell)$,

$$\alpha_\ell^* = \frac{\mathbb{C}[\hat{T}_\ell^{(\ell)}, \hat{U}_\ell^{(\ell)}]}{\mathbb{V}[\hat{U}_\ell^{(\ell)}]}, \quad (3.12)$$

and the resulting variance of the control variate estimator of the correction is

$$\mathbb{V}[\hat{T}_\ell^{\text{CV}}(\alpha_\ell^*)] = (1 - \rho_\ell^2)\mathbb{V}[\hat{T}_\ell^{(\ell)}], \quad \text{with } \rho_\ell = \frac{\mathbb{C}[\hat{T}_\ell^{(\ell)}, \hat{U}_\ell^{(\ell)}]}{\mathbb{V}[\hat{T}_\ell^{(\ell)}]^{1/2}\mathbb{V}[\hat{U}_\ell^{(\ell)}]^{1/2}} \in [-1, 1]. \quad (3.13)$$

The correction estimators $(\hat{T}_\ell^{\text{CV}})_{\ell=0}^L$ being mutually independent, the variance of the optimal MLMC-CV estimator is

$$\mathbb{V}[\hat{\theta}^{\text{MLMC-CV}}(\alpha_0^*, \dots, \alpha_L^*)] = \sum_{\ell=0}^L \mathbb{V}[\hat{T}_\ell^{\text{CV}}(\alpha_\ell^*)] = \sum_{\ell=0}^L (1 - \rho_\ell^2)\mathbb{V}[\hat{T}_\ell^{(\ell)}] \leq \sum_{\ell=0}^L \mathbb{V}[\hat{T}_\ell^{(\ell)}], \quad (3.14)$$

indicating that the variance of the MLMC-CV estimator is smaller (as long as $\rho_\ell^2 > 0$) than that of the MLMC estimator; see eq. (3.5). We remark that the variance reduction depends on the (squared) correlation between $\hat{T}_\ell^{(\ell)}$ and $\hat{U}_\ell^{(\ell)}$, which, in turn, typically relates to some measure of similarity between high-fidelity corrections $Y_\ell - Y_{\ell-1}$ and the corresponding CV corrections (see appendix A for the expectation and variance estimators in a single-level setting).

In our surrogate-based approach, we may define control variates as $Z_\ell = g_\ell(\mathbf{X})$, where g_ℓ is a surrogate of f_ℓ , so that τ_ℓ and $\tau_{\ell-1}$ could be estimated using samples of Z_ℓ and $Z_{\ell-1}$, respectively. It is then crucial to construct the surrogates such that their successive differences $g_\ell - g_{\ell-1}$ are good approximations of $f_\ell - f_{\ell-1}$, to ensure a high similarity between $Y_\ell - Y_{\ell-1}$ and $Z_\ell - Z_{\ell-1}$. However,

constructing g_ℓ as a surrogate of f_ℓ does not give any guarantee on the quality of $g_\ell - g_{\ell-1}$ as a surrogate of $f_\ell - f_{\ell-1}$. Instead, in addition to the surrogates models g_ℓ of f_ℓ , we construct surrogate models h_ℓ of the differences $f_\ell - f_{\ell-1}$, and we define auxiliary surrogate models $\tilde{g}_{\ell-1} = g_\ell - h_\ell$, for $\ell = 1, \dots, L$. On each level ℓ , we then use samples of $Z_\ell = g_\ell(\mathbf{X})$ and $\tilde{Z}_{\ell-1} = \tilde{g}_{\ell-1}(\mathbf{X})$ for the estimation of τ_ℓ and $\tau_{\ell-1}$, respectively. As a result, the variance reduction now depends on the similarity between $Y_\ell - Y_{\ell-1}$ and $W_\ell := Z_\ell - \tilde{Z}_{\ell-1} = g_\ell(\mathbf{X}) - \tilde{g}_{\ell-1}(\mathbf{X}) = h_\ell(\mathbf{X})$, where h_ℓ has been constructed to ensure such a similarity. Specifically, for the expectation, the MLMC-CV estimator reads

$$\hat{E}^{\text{MLMC-CV}}[Y](\alpha_0, \dots, \alpha_L) = \sum_{\ell=0}^L \left(\hat{E}^{(\ell)}[Y_\ell] - \hat{E}^{(\ell)}[Y_{\ell-1}] \right) - \alpha_\ell \left(\hat{E}^{(\ell)}[Z_\ell] - \hat{E}^{(\ell)}[\tilde{Z}_{\ell-1}] - (\mu_{Z_\ell} - \mu_{\tilde{Z}_{\ell-1}}) \right), \quad (3.15)$$

with optimal values of α_ℓ given by (see appendix A.1, with $M = 1$)

$$\alpha_\ell^* = \frac{\mathbb{C}[Y_\ell - Y_{\ell-1}, W_\ell]}{\mathbb{V}[Z_\ell - \tilde{Z}_{\ell-1}]}, \quad \forall \ell = 0, \dots, L, \quad (3.16)$$

resulting in level-dependent reduction factors $1 - \rho_\ell^2$, where

$$\rho_\ell^2 = \frac{\mathbb{C}[Y_\ell - Y_{\ell-1}, W_\ell]^2}{\mathbb{V}[Y_\ell - Y_{\ell-1}]\mathbb{V}[W_\ell]} \quad (3.17)$$

is the squared correlation coefficient between $Y_\ell - Y_{\ell-1}$ and $W_\ell = h_\ell(\mathbf{X})$. It should be noted that, because of the linearity of the expectation operator and its MC estimator, the use of $\tilde{g}_{\ell-1}$ is superfluous. Indeed, we may directly define the MLMC-CV estimator of the expectation as

$$\hat{E}^{\text{MLMC-CV}}[Y](\alpha_0, \dots, \alpha_L) = \sum_{\ell=0}^L \left(\hat{E}^{(\ell)}[Y_\ell] - \hat{E}^{(\ell)}[Y_{\ell-1}] \right) - \alpha_\ell \left(\hat{E}^{(\ell)}[W_\ell] - \mu_{W_\ell} \right), \quad (3.18)$$

with $\mu_{W_\ell} = \mathbb{E}[W_\ell]$. For the variance, the MLMC-CV estimator reads

$$\hat{V}^{\text{MLMC-CV}}[Y](\alpha_0, \dots, \alpha_L) = \sum_{\ell=0}^L \left(\hat{V}^{(\ell)}[Y_\ell] - \hat{V}^{(\ell)}[Y_{\ell-1}] \right) - \alpha_\ell \left(\hat{V}^{(\ell)}[Z_\ell] - \hat{V}^{(\ell)}[\tilde{Z}_{\ell-1}] - (\sigma_{Z_\ell}^2 - \sigma_{\tilde{Z}_{\ell-1}}^2) \right), \quad (3.19)$$

with $\sigma_{Z_\ell}^2 = \mathbb{V}[Z_\ell]$ and $\sigma_{\tilde{Z}_{\ell-1}}^2 = \mathbb{V}[\tilde{Z}_{\ell-1}]$. The resulting level-dependent reduction factors are then related to the correlation between $(Y_\ell - Y_{\ell-1} - \mathbb{E}[Y_\ell - Y_{\ell-1}])^2$ and $(h_\ell(\mathbf{X}) - \mathbb{E}[h_\ell(\mathbf{X})])^2$ (see appendix A.2, with $M = 1$). Further details on the construction of h_ℓ will be given in section 4.

3.2.3 Multilevel Monte Carlo with multilevel control variates (MLMC-MLCV)

We now propose to further improve the MLMC-CV estimator eq. (3.10) by combining MLMC with the MLCV approach described in section 3.2.1, resulting in the MLMC-MLCV method. The approach consists in using, at each level ℓ , the surrogate-based control variates of *all the levels* $\ell' = 0, \dots, L$, rather than only using those of level ℓ , as was previously done with the MLMC-CV approach.

The MLMC-MLCV estimator then reads

$$\hat{\theta}^{\text{MLMC-MLCV}}(\boldsymbol{\alpha}_0, \dots, \boldsymbol{\alpha}_L) = \sum_{\ell=0}^L \hat{T}_\ell^{\text{MLCV}}(\boldsymbol{\alpha}_\ell), \quad (3.20)$$

where $\boldsymbol{\alpha}_\ell$ denotes the CV parameter at level ℓ , and $\hat{T}_\ell^{\text{MLCV}}(\boldsymbol{\alpha}_\ell)$ is the MLCV estimator of T_ℓ ,

$$\hat{T}_\ell^{\text{MLCV}}(\boldsymbol{\alpha}_\ell) = \hat{T}_\ell^{(\ell)} - \boldsymbol{\alpha}_\ell^\top (\hat{\mathbf{U}}_\ell^{(\ell)} - \mathbf{U}_\ell), \quad (3.21)$$

with

$$\mathbf{U}_0 = (\tau_k)_{k=0}^L, \quad \hat{\mathbf{U}}_0^{(0)} = (\hat{\tau}_k^{(0)})_{k=0}^L \quad (3.22)$$

$$\mathbf{U}_\ell = (U_k)_{k=1}^L, \quad \text{for } \ell > 0, \quad \hat{\mathbf{U}}_\ell^{(\ell)} = (\hat{U}_k^{(\ell)})_{k=1}^L, \quad \text{for } \ell > 0, \quad (3.23)$$

and with U_k and $\hat{U}_k^{(\ell)}$ defined as in section 3.2.2.

Because each term eq. (3.21) is an unbiased (multiple) CV estimator of T_ℓ , the resulting estimator eq. (3.20) is also unbiased, and the optimal (variance minimizing) value $\boldsymbol{\alpha}_\ell^*$ of $\boldsymbol{\alpha}_\ell$ is given individually for each $\ell = 0, \dots, L$ as the optimal (multiple) CV parameter for $\hat{T}_\ell^{(\ell)}(\boldsymbol{\alpha}_\ell)$,

$$\boldsymbol{\alpha}_\ell^* = \mathbb{C}[\hat{\mathbf{U}}_\ell^{(\ell)}]^{-1} \mathbb{C}[\hat{\mathbf{U}}_\ell^{(\ell)}, \hat{T}_\ell^{(\ell)}]. \quad (3.24)$$

The resulting variance is given by

$$\mathbb{V}[\hat{\theta}^{\text{MLMC-MLCV}}(\boldsymbol{\alpha}_0^*, \dots, \boldsymbol{\alpha}_L^*)] = \sum_{\ell=0}^L \mathbb{V}[\hat{T}_\ell^{\text{MLCV}}(\boldsymbol{\alpha}_\ell^*)] = \sum_{\ell=0}^L (1 - R_\ell^2) \mathbb{V}[\hat{T}_\ell^{(\ell)}] \quad (3.25)$$

with

$$R_\ell^2 = \mathbb{V}[\hat{T}_\ell^{(\ell)}]^{-1} \mathbb{C}[\hat{\mathbf{U}}_\ell^{(\ell)}, \hat{T}_\ell^{(\ell)}]^\top \boldsymbol{\alpha}_\ell^* \in [0, 1]. \quad (3.26)$$

Again, owing to the fact that $R_\ell^2 \leq 1$, the variance of the MLMC-MLCV estimator is always less than or equal to the variance of the MLMC estimator given by $\mathbb{V}[\hat{\theta}_L^{\text{MLMC}}] = \sum_{\ell=0}^L \mathbb{V}[\hat{T}_\ell^{(\ell)}]$ (see eq. (3.5)).

In our surrogate-based approach, $L + 1$ surrogate-based control variates can be used for the coarsest level $\ell = 0$, namely g_0, \dots, g_L , so that $\boldsymbol{\alpha}_0 \in \mathbb{R}^{L+1}$. At correction levels $\ell > 0$, we can use L control variates based on g_1, \dots, g_L and $\tilde{g}_0, \dots, \tilde{g}_{L-1}$, as described in section 3.2.2, so that $\boldsymbol{\alpha}_\ell \in \mathbb{R}^L$, for $\ell > 0$.

The MLMC-MLCV estimator of the expectation $\mathbb{E}[f(\mathbf{X})]$ is defined by eq. (3.20), with

$$\hat{T}_0^{\text{MLCV}}(\boldsymbol{\alpha}_0) = \hat{E}^{(0)}[Y_0] - \boldsymbol{\alpha}_0^\top (\hat{E}^{(0)}[\mathbf{Z}] - \boldsymbol{\mu}_\mathbf{Z}), \quad (3.27)$$

$$\hat{T}_\ell^{\text{MLCV}}(\boldsymbol{\alpha}_\ell) = \hat{E}^{(\ell)}[Y_\ell] - \hat{E}^{(\ell)}[Y_{\ell-1}] - \boldsymbol{\alpha}_\ell^\top (\hat{E}^{(\ell)}[\mathbf{W}] - \boldsymbol{\mu}_\mathbf{W}), \quad \text{for } \ell > 0, \quad (3.28)$$

where

$$\mathbf{Z} = (Z_0, \dots, Z_L) = (g_0(\mathbf{X}), \dots, g_L(\mathbf{X})), \quad \boldsymbol{\mu}_\mathbf{Z} = \mathbb{E}[\mathbf{Z}], \quad (3.29)$$

$$\mathbf{W} = (W_1, \dots, W_L) = (h_1(\mathbf{X}), \dots, h_L(\mathbf{X})), \quad \boldsymbol{\mu}_\mathbf{W} = \mathbb{E}[\mathbf{W}]. \quad (3.30)$$

The optimal values $\boldsymbol{\alpha}_\ell^*$ of the CV parameters are given by

$$\boldsymbol{\alpha}_0^* = \boldsymbol{\Sigma}_0^{-1} \mathbf{c}_0, \quad \boldsymbol{\Sigma}_0 = \mathbb{C}[\mathbf{Z}], \quad \mathbf{c}_0 = \mathbb{C}[\mathbf{Z}, Y_0], \quad (3.31)$$

$$\boldsymbol{\alpha}_\ell^* = \boldsymbol{\Sigma}_\ell^{-1} \mathbf{c}_\ell, \quad \boldsymbol{\Sigma}_\ell = \mathbb{C}[\mathbf{W}], \quad \mathbf{c}_\ell = \mathbb{C}[\mathbf{W}, Y_\ell - Y_{\ell-1}], \quad \text{for } \ell > 0, \quad (3.32)$$

resulting in $R_\ell^2 = \mathbb{V}[Y_\ell - Y_{\ell-1}]^{-1} \mathbf{c}_\ell^\top \boldsymbol{\Sigma}_\ell \mathbf{c}_\ell$, for $\ell = 0, \dots, L$. Note that $\boldsymbol{\Sigma}_\ell = \mathbb{C}[\mathbf{W}]$ is the same for all $\ell > 0$. The optimal MLMC-MLCV variance estimator is derived in appendix C.

Remark 3.1. In practice, Σ and \mathbf{c}_ℓ may be estimated using either a pilot sample or the same sample as for the estimation of $\hat{\mathbf{U}}_\ell^{(\ell)}$ and $\hat{T}_\ell^{(\ell)}$. Alternatively, in the specific context of PC-based control variates, for the estimation of the expectation, a closed-form expression for Σ can be obtained. Letting $Z_\ell = g_\ell(\mathbf{X}) = \sum_{k=0}^{P_g^\ell} \mathfrak{g}_{\ell,k} \Psi_k(\mathbf{X})$ and $W_\ell = h_\ell(\mathbf{X}) = \sum_{k=0}^{P_h^\ell} \mathfrak{h}_{\ell,k} \Psi_k(\mathbf{X})$, we have

$$\forall m, m' = 0, \dots, L, \quad [\Sigma_0]_{m,m'} = \mathbb{C}[Z_m, Z_{m'}] = \sum_{k=1}^{\min(P_g^m, P_g^{m'})} \mathfrak{g}_{m,k} \mathfrak{g}_{m',k}, \quad (3.33)$$

$$\forall m, m' = 1, \dots, L, \quad [\Sigma_\ell]_{m,m'} = \mathbb{C}[W_m, W_{m'}] = \sum_{k=1}^{\min(P_h^m, P_h^{m'})} \mathfrak{h}_{m,k} \mathfrak{h}_{m',k}, \quad \text{for } \ell > 0. \quad (3.34)$$

3.3 Practical details

A summary of the methods is presented in table 1, including two minor variations of the MLMC-CV and MLMC-MLCV estimators introduced previously, which consist of using fewer surrogate models for the multilevel estimation. We remark that all the MLMC-like estimators (including the MLMC estimator), hereafter abbreviated MLMC-*, have variance $\sum_{\ell=0}^L (1 - R_\ell^2) \mathbb{V}[\hat{T}_\ell^{(\ell)}]$, with

- $R_\ell^2 = 0$ for plain MLMC;
- $R_\ell^2 = \rho_\ell^2$ as defined in eq. (3.13) for MLMC-CV; and
- R_ℓ^2 defined by eq. (3.26) for MLMC-MLCV.

For all these methods, we will further assume that $\mathbb{V}[\hat{T}_\ell^{(\ell)}] \lesssim n_\ell^{-1} \mathcal{V}_\ell$ (see section 3.1), which implies that $\mathbb{V}[\hat{T}_\ell^{\text{MLMC-*}}(\boldsymbol{\alpha}_\ell^*)] = (1 - R_\ell^2) \mathbb{V}[\hat{T}_\ell^{(\ell)}] \lesssim n_\ell^{-1} \mathcal{V}_\ell^{\text{CV}}$, with $\mathcal{V}_\ell^{\text{CV}} := (1 - R_\ell^2) \mathcal{V}_\ell$.

In the surrogate-based variants of MLMC, the cost of evaluating the surrogate models is assumed to be negligible compared to the costs of evaluating the simulators f_0, \dots, f_L . Therefore, the total computational cost of the MLMC-* estimator reduces to the cost of the MLMC estimator, $\text{cost}(\hat{\theta}^{\text{MLMC-MLCV}}) = \text{cost}(\hat{\theta}^{\text{MLMC}})$, given by eq. (3.6). Similarly to the case of the MLMC estimator (see, e.g., [36]), the optimal sample sizes $(n_\ell^*)_{\ell=0}^L$ such that $\sum_{\ell=0}^L \mathcal{V}_\ell^{\text{CV}} / n_\ell^*$ is minimal under a constrained computational budget of \mathcal{C} are given by

$$n_\ell^* = \frac{\mathcal{C}}{\mathcal{S}_L} \sqrt{\frac{\mathcal{V}_\ell^{\text{CV}}}{\mathcal{C}_\ell + \mathcal{C}_{\ell-1}}}, \quad \text{with } \mathcal{S}_\ell := \sum_{\ell'=0}^{\ell} \sqrt{(\mathcal{C}_{\ell'} + \mathcal{C}_{\ell'-1}) \mathcal{V}_{\ell'}^{\text{CV}}}, \quad (3.35)$$

so that $\sum_{\ell=0}^L \mathcal{V}_\ell^{\text{CV}} / n_\ell^* = \mathcal{S}_L^2 / \mathcal{C}$. As a consequence,

$$\mathbb{V}[\hat{\theta}_L^{\text{MLMC-MLCV}}(\boldsymbol{\alpha}_L^*)] \lesssim \frac{\mathcal{S}_L^2}{\mathcal{C}}, \quad (3.36)$$

with an equality between the left- and right-hand sides for the expectation estimators.

In practice, $\mathcal{V}_\ell^{\text{CV}}$ is not known and must be estimated for each level. In this work, we consider the sequential algorithm proposed in [36, Algorithm 2] for the MLMC. The algorithm starts from an initial, small number of samples n_ℓ^{init} on each level. Then, it selects the optimal level on which to increase the sample size by an inflation factor $r_\ell > 1$, i.e. the level ℓ^* on which the reduction in total variance relative to the additional computational effort achieved by inflating the sample size by r_{ℓ^*} is maximal.

Table 1: Summary of the methods. The first three are state-of-the-art methods, the next three are the novel multilevel methods proposed in this paper, and the last two are variants of the proposed methods.

Method	Form of the estimator	Eq.
Monte Carlo (MC) .	$\hat{\theta}$	
Control Variates (CV) [29, 30, 31, 38, 33].	$\hat{\theta} - \sum_{m=1}^M \alpha_m (\hat{\tau}_m - \tau_m)$.	2.4
Multilevel Monte Carlo (MLMC) [21, 22].	$\hat{\theta}_0^{(0)} + \sum_{\ell=1}^L \hat{T}_\ell^{(\ell)}$.	3.2
Multilevel Control Variates (MLCV) . CV method where the CVs are based on surrogate models of simulators of different levels of fidelity.	$\hat{\theta} - \sum_{\ell=0}^L \alpha_\ell (\hat{\tau}_\ell - \tau_\ell)$.	2.4
MLMC-CV . MLMC with one CV at each correction level based on surrogate models of the simulators on the corresponding level. Corresponds to a surrogate-based MLMF [18, 19] where the exact statistics of the CVs are known.	$\hat{\theta}_0^{(0)} - \alpha_0 (\hat{\tau}_0^{(0)} - \tau_0) + \sum_{\ell=1}^L \left(\hat{T}_\ell^{(\ell)} - \alpha_\ell (\hat{U}_\ell^{(\ell)} - U_\ell) \right)$.	3.10, 3.11
MLMC-MLCV . MLMC-CV with the CVs based on the surrogate models g_0, \dots, g_L on level 0 and on h_1, \dots, h_L on levels $\ell > 0$.	$\hat{\theta}_0^{(0)} - \alpha_0^{\mathbf{I}} \begin{bmatrix} \hat{\tau}_0^{(0)} - \tau_0 \\ \vdots \\ \hat{\tau}_L^{(0)} - \tau_L \end{bmatrix} + \sum_{\ell=1}^L \left(\hat{T}_\ell^{(\ell)} - \alpha_\ell^{\mathbf{I}} \begin{bmatrix} \hat{U}_1^{(\ell)} - U_1 \\ \vdots \\ \hat{U}_L^{(\ell)} - U_L \end{bmatrix} \right)$.	3.20, 3.21
MLMC-CV[0] . MLMC-CV using only one CV based on the surrogate g_0 on level 0 and no CVs on levels $\ell > 0$	$\hat{\theta}_0^{(0)} - \alpha_0 (\hat{\tau}_0^{(0)} - \tau_0) + \sum_{\ell=1}^L \hat{T}_\ell^{(\ell)}$.	
MLMC-MLCV[0] . MLMC-MLCV using only CVs based on the surrogates g_0 and g_1 on level 0, and h_1 on levels $\ell > 0$.	$\hat{\theta}_0^{(0)} - \alpha_0^{\mathbf{I}} \begin{bmatrix} \hat{\tau}_0^{(0)} - \tau_0 \\ \hat{\tau}_1^{(0)} - \tau_1 \end{bmatrix} + \sum_{\ell=1}^L \left(\hat{T}_\ell^{(\ell)} - \alpha_\ell (\hat{U}_1^{(\ell)} - U_1) \right)$.	

Algorithm 1 Simplified MLMC-* algorithm inspired by [36].

Require: $n_\ell^{\text{init}} > 1$, $r_\ell > 1$, surrogate models (depending on the method), and budget \mathcal{C} .

- 1: Set consumed budget to $\tilde{\mathcal{C}} = 0$ and $\delta n_\ell = n_\ell^{\text{init}}$ samples on levels $\ell \leq L$;
 - 2: **while** $\tilde{\mathcal{C}} \leq \mathcal{C}$ **do**
 - 3: compute δn_ℓ samples on each level by evaluating f_ℓ and the appropriate surrogates;
 - 4: update sample size on each level: $n_\ell \leftarrow n_\ell + \delta n_\ell$;
 - 5: update consumed budget: $\tilde{\mathcal{C}} \leftarrow \tilde{\mathcal{C}} + \sum_{\ell=0}^L \delta n_\ell (\mathcal{C}_\ell + \mathcal{C}_{\ell-1})$;
 - 6: estimate the optimal CV parameter(s) on each level;
 - 7: compute/update CV estimates for $\hat{T}_\ell^{(\ell)}$ and $\mathcal{V}_\ell^{\text{CV}}$ from samples on levels $\ell \leq L$;
 - 8: select level $\ell^* = \arg \max_{0 \leq \ell \leq L} \frac{\mathcal{V}_\ell^{\text{CV}}}{r_\ell n_\ell^2 (\mathcal{C}_\ell + \mathcal{C}_{\ell-1})}$;
 - 9: $\delta n_{\ell^*} \leftarrow \lfloor (r_{\ell^*} - 1)n_{\ell^*} \rfloor$, $\delta n_{\ell \neq \ell^*} \leftarrow 0$;
 - 10: **end while**
 - 11: **return** $\hat{\theta}_L^{\text{MLMC-*}}$, the MLMC-* estimate of θ_L .
-

4 Numerical experiments

We demonstrate the value of our MLMC-MLCV method on the uncertain heat equation problem proposed in [18], with seven random input variables modeling the uncertainty in the diffusion coefficient and the initial condition. Although academic, this problem is nonetheless challenging because of its moderate stochastic dimensionality, the non-linearity of the solution with respect to the inputs, and the high variance of the random output. The test problem is summarized in section 4.1. The surrogate models used for the control variates are described in section 4.2, and the results from numerical experiments are presented and discussed in section 4.3.

4.1 Problem description

We consider the partial differential equation describing the time-evolution of the temperature $u(x, t; \mathbf{X})$ in a 1D rod of unit length over the time interval $[0, T]$, with uncertain (random) initial data u_0 and thermal diffusivity ν ,

$$\begin{cases} \frac{\partial u(x, t; \mathbf{X})}{\partial t} = \nu(\mathbf{X}) \frac{\partial^2 u(x, t; \mathbf{X})}{\partial x^2}, & x \in \mathcal{D} := (0, 1), \quad t \in [0, T], \\ u(x, 0; \mathbf{X}) = u_0(x; \mathbf{X}), & x \in \mathcal{D}, \\ u(0, t; \mathbf{X}) = u(1, t; \mathbf{X}) = 0, & t \in [0, T], \end{cases} \quad (4.1)$$

where $\mathbf{X}: \Omega \rightarrow \Xi$ is a random vector modelling the uncertainty in the input parameters, and where $\nu(\mathbf{X}) > 0$ almost surely. The solution of eq. (4.1) may be expressed as

$$u(x, t; \mathbf{X}) = \sum_{k=1}^{\infty} a_k(\mathbf{X}) \exp(-\nu(\mathbf{X})k^2\pi^2t) \sin(k\pi x) \quad (4.2)$$

with

$$a_k(\mathbf{X}) = 2 \int_{\mathcal{D}} u_0(x; \mathbf{X}) \sin(k\pi x) dx. \quad (4.3)$$

The initial condition is chosen to have the same prescribed form as in [18]. Specifically, we consider $u_0(x; \mathbf{X}) = \mathcal{G}(\mathbf{X})\mathcal{F}_1(x) + \mathcal{I}(\mathbf{X})\mathcal{F}_2(x)$ with

$$\mathcal{F}_1(x) = \sin(\pi x), \quad (4.4)$$

$$\mathcal{F}_2(x) = \sin(2\pi x) + \sin(3\pi x) + 50(\sin(9\pi x) + \sin(21\pi x)), \quad (4.5)$$

$$\mathcal{I}(\mathbf{X}) = \frac{7}{2} [\sin(X_1) + 7\sin(X_2)^2 + 0.1X_3^4 \sin(X_1)], \quad (4.6)$$

$$\mathcal{G}(\mathbf{X}) = 50(4|X_5| - 1)(4|X_6| - 1)(4|X_7| - 1), \quad (4.7)$$

which allows to control the spectral content of the solution u . Furthermore, as in [18], the diffusion coefficient is modelled by $\nu(\mathbf{X}) = X_4$. The random output variable of interest is defined as the integral of the temperature along the rod at final time T ,

$$\mathcal{M}(\mathbf{X}) = \int_{\mathcal{D}} u(x, T; \mathbf{X}) dx \quad (4.8)$$

$$= \sum_{k=1}^{\infty} a_k(\mathbf{X}) \int_{\mathcal{D}} \exp(-\nu(\mathbf{X})k^2\pi^2T) \sin(k\pi x) dx \quad (4.9)$$

$$= \mathcal{G}(\mathbf{X})\mathcal{H}_1(\mathbf{X}) + \mathcal{I}(\mathbf{X}) [\mathcal{H}_3(\mathbf{X}) + 50\mathcal{H}_9(\mathbf{X}) + 50\mathcal{H}_{21}(\mathbf{X})], \quad (4.10)$$

where $\mathcal{H}_k(\mathbf{X}) = \frac{2}{k\pi} \exp(-\nu(\mathbf{X})k^2\pi^2T)$. In this experiment, we seek to estimate the expectation $\mathbb{E}[\mathcal{M}(\mathbf{X})]$, for a given uncertain setting. Consistently with [18], we consider the random variables X_1, \dots, X_7 to be independent and distributed as

$$X_1, X_2, X_3 \sim \mathcal{U}[-\pi, \pi], \quad X_4 \sim \mathcal{U}[\nu_{\min}, \nu_{\max}], \quad X_5, X_6, X_7 \sim \mathcal{U}[-1, 1]. \quad (4.11)$$

The expected value $\mathbb{E}[\mathcal{M}(\mathbf{X})]$ is then given by

$$\mathbb{E}[\mathcal{M}(\mathbf{X})] = 50H_1 + \frac{49}{4} (H_3 + 50H_9 + 50H_{21}), \quad (4.12)$$

where

$$H_k = \mathbb{E}[\mathcal{H}_k(\mathbf{X})] = \frac{2}{k^3\pi^3T} \frac{\exp(-\nu_{\min}k^2\pi^2T) - \exp(-\nu_{\max}k^2\pi^2T)}{\nu_{\max} - \nu_{\min}}. \quad (4.13)$$

Finally, we set $T = 0.5$, $\nu_{\min} = 0.001$ and $\nu_{\max} = 0.009$, resulting in $\mathbb{E}[\mathcal{M}(\mathbf{X})] \approx 41.98$.

Numerically, $\mathcal{M}(\mathbf{X})$ is approximated by truncating the Fourier expansion in eq. (4.9) to $K < \infty$ modes and by approximating the integrals in eqs. (4.3) and (4.9) by a trapezoidal quadrature rule with equispaced nodes in $[0, 1]$. The multilevel hierarchy of simulators $\{f_\ell\}_{\ell=0}^L$ is then defined according to the number of quadrature nodes N_ℓ used for the approximation at level ℓ . Specifically, $\mathcal{M}(\mathbf{X})$ is approximated at level ℓ by

$$Y_\ell = f_\ell(\mathbf{X}) = \sum_{k=1}^K A_k^\ell(\mathbf{X}) B_k^\ell(\mathbf{X}), \quad (4.14)$$

with

$$A_k^\ell(\mathbf{X}) = 2 \sum_{i=1}^{N_\ell} w_i u_0(x_i; \mathbf{X}) \sin(k\pi x_i), \quad B_k^\ell(\mathbf{X}) = \exp(-\nu(\mathbf{X})k^2\pi^2T) \sum_{i=1}^{N_\ell} w_i \sin(k\pi x_i), \quad (4.15)$$

where $\{(x_i, w_i)\}_{i=1}^{N_\ell}$ are the pairs of quadrature nodes and associated weights on level ℓ . It is then natural to assume that the computational cost \mathcal{C}_ℓ of an evaluation of f_ℓ is $\mathcal{O}(KN_\ell)$. The statistic

of interest is thus $\theta = \theta_L = \mathbb{E}[f_L(\mathbf{X})]$, whose MC estimator $\hat{\theta}_L = n_L^{-1} \sum_{i=1}^{n_L} f_L(\mathbf{X}^{(L,i)})$ will represent the baseline estimator for our experiments. Besides, the quality of all the presented estimators will be assessed in terms of their root mean square error (RMSE) w.r.t. the exact statistic $\mathbb{E}[\mathcal{M}(\mathbf{X})]$ given by eq. (4.12).

In the following experiments, we set the number of quadrature nodes $K = 21$ and the number of levels to 4 (i.e. $L = 3$). Furthermore, we set $N_\ell = 120 \times 2^{L-\ell}$ so that evaluating f_ℓ is twice as expensive as evaluating $f_{\ell-1}$. Table 2 summarizes the number of quadrature nodes and the evaluation cost per level. Note that the costs are normalized so that $\mathcal{C}_3 = 1$.

Table 2: Number of quadrature nodes N_ℓ , simulator evaluation cost \mathcal{C}_ℓ and MLMC correction evaluation cost $\mathcal{C}_\ell + \mathcal{C}_{\ell-1}$ per level.

ℓ	0	1	2	3
N_ℓ	15	30	60	120
\mathcal{C}_ℓ	0.125	0.25	0.5	1
$\mathcal{C}_\ell + \mathcal{C}_{\ell-1}$	0.125	0.375	0.75	1.5

4.2 Surrogate models

We will mostly use PC models (see section 2.4.3) for the surrogate-based CV estimators. Constructing high-quality surrogates in 7 dimensions can be hard, especially in the presence of non-linearities and with a limited sample size. To avoid overfitting, we resort to the *least angle regression* (LARS) procedure [14], which is a model-selection regression method that promotes sparsity. More precisely, we employ the basis-adaptive hybrid LARS algorithm proposed by [7, Fig. 5] for the selection of sparse PC bases. For a given design of experiment (DoE) for the PC surrogate construction, this algorithm applies the LARS procedure on candidate PC bases \mathcal{A}_p of increasing total polynomial degree $p = 1, \dots, p_{\max}$, resulting for each p in the selection of a limited number $|\tilde{\mathcal{A}}_p| < |\mathcal{A}_p|$ of basis polynomial functions. For each p , a PC surrogate is constructed by classical least-squares regression on the corresponding reduced (or active) PC basis $\tilde{\mathcal{A}}_p$, and its quality is estimated using a corrected leave-one-out cross-validation procedure [7]. The best surrogate according to this quality measure is eventually retained, and the associated reduced PC basis is denoted by $\tilde{\mathcal{A}}_{p^*}$.

In our experiments, we set $p_{\max} = 16$ and the construction budget to 400 times the evaluation cost of f_3 , i.e. $\mathcal{C}^{\text{DoE}} = 400$. This budget is distributed equally among the different levels, so that the associated evaluation cost on each level corresponds to the cost of 100 f_3 evaluations, i.e. $n_\ell^{\text{DoE}} \mathcal{C}_\ell = 100$, as reported in tables 3 and 4. Once constructed, the quality of a surrogate g of f is assessed in terms of the Q^2 measure,

$$Q^2(g, f) = 1 - \frac{\mathbb{E}[(g(\mathbf{X}) - f(\mathbf{X}))^2]}{\mathbb{V}[f(\mathbf{X})]}, \quad (4.16)$$

estimated using a test sample of size $n_{\text{test}} = 10\,000$, which is more robust than the corrected leave-one-out measure used for the model selection. The higher the Q^2 value, the higher-quality the associated surrogate.

For the MLCV estimator, we learn a PC model g_ℓ of f_ℓ from a training DoE of size n_ℓ^{DoE} generated by latin hypercube sampling (LHS) improved by simulated annealing. The training DoE samples for the surrogate construction are different for each level and independent. Table 3 summarizes the properties of the different PC models. Except for g_3 , which is built with only 100 points, all the PC

models have a good Q^2 , greater than 0.8. We observe that the basis-adaptive LARS algorithm has selected a decreasing polynomial degree p^* with level ℓ , while retaining only a limited number $|\tilde{\mathcal{A}}_{p^*}|$ of polynomials in these reduced bases. Thus, although g_1 and g_2 have distinct degrees, the sizes of the associated reduced bases are similar.

Table 3: PC models for CV and MLCV estimators, with their sample size, degree and quality measure. These models are built with the basis-adaptive LARS algorithm of [7, Fig. 5].

PC models	g_0	g_1	g_2	g_3
n_ℓ^{DoE}	800	400	200	100
p^*	10	8	6	4
$ \tilde{\mathcal{A}}_{p^*} $	211	77	74	14
Q^2	0.98	0.94	0.86	0.59

For the MLMC-based estimators, we learn a PC model g_ℓ of f_ℓ for $\ell = 0, \dots, L$ and a PC model h_ℓ of $f_\ell - f_{\ell-1}$ for $\ell = 1, \dots, L$. In practice, training points used for the construction of g_0, \dots, g_L may be reused for the construction of h_1, \dots, h_L using nested DoEs. However, it should be noted that the generation of nested LHS DoEs is not as straightforward as for purely random DoEs. While using nested random DoEs is a perfectly valid strategy, we opt for an alternative choice based on LHS. Specifically, we first generate a DoE $\mathcal{X}_{\text{DoE}}^{(0)}$ of size n_0^{DoE} using LHS improved by simulated annealing. Then, for $\ell = 1, \dots, L$, we sequentially extract a DoE $\mathcal{X}_{\text{DoE}}^{(\ell)} \subset \mathcal{X}_{\text{DoE}}^{(\ell-1)}$ of size n_ℓ^{DoE} . The subset $\mathcal{X}_{\text{DoE}}^{(\ell)}$ is selected such that it has minimal centered L^2 discrepancy among a pool of 1 million random candidate subsets. Table 4 summarizes the properties of these different PC models. The PC models again have a quality measure higher than 0.8, except for g_3 and h_3 , which were constructed from only 100 evaluations. Note that the values of p^* , $|\tilde{\mathcal{A}}_{p^*}|$ and Q^2 are of the same order of magnitude as those of the PC surrogate models reported in table 3.

Table 4: PC models for the estimators combining MLMC and CV techniques, with their sample size, degree and quality measure. More precisely, g_ℓ is the surrogate of f_ℓ , h_ℓ is the surrogate of $f_\ell - f_{\ell-1}$ and the samples used to train h_ℓ are also used to train g_ℓ and $g_{\ell-1}$. These models are built with the basis-adaptive LARS algorithm of [7, Fig. 5].

PC models	g_0	g_1	g_2	g_3	h_1	h_2	h_3
n_ℓ^{DoE}	800	400	200	100	400	200	100
p^*	10	8	8	2	9	7	5
$ \tilde{\mathcal{A}}_{p^*} $	211	83	57	6	112	71	30
Q^2	0.98	0.95	0.86	0.21	0.99	0.82	0.63

Hereafter, the CV and MLCV estimators use the PC models of table 3, while the MLMC-based estimators use the PC models presented in table 4.

4.3 Results

First, in section 4.3.1, we illustrate the use of one or several control variates to reduce the variance of a single-level MC estimator. Then, we compare the MLCV and MLMC-MLCV approaches with the MC and MLMC estimators in section 4.3.2, and we discuss variants of the MLMC-CV and MLMC-MLCV approaches, considering only a limited subset of the surrogate models, in section 4.3.3. We

conclude the analysis by reporting the estimation budget allocation across levels resulting from the various MLMC-based methods in section 4.3.4. In practice, all the MLMC-based estimators are built using algorithm 1. Unless stated otherwise, the parameters are set to $n_\ell^{\text{init}} = 30$ and $r_\ell = 1.1$ for $\ell = 0, \dots, 3$. The quality of the various estimators will be assessed in terms of their RMSE w.r.t. $\mathbb{E}[\mathcal{M}(\mathbf{X})]$, estimated from 500 repetitions of the experiment.

4.3.1 Single-level MC and control variates


In this first part, we consider only the finest simulator f_3 and try to reduce the variance of the MC estimator of $\mathbb{E}[f_3(\mathbf{X})]$ by means of surrogate-based CVs. For that purpose, we consider a first-order Taylor polynomial expansion $g_3^{T_1}$ around $\boldsymbol{\mu}_{\mathbf{X}} = \mathbb{E}[\mathbf{X}]$ (see section 2.4.2 and appendix D) and the PC model g_3^{PC} described in table 3.

Table 5: Pearson correlation coefficients between the finest simulator f_3 and the corresponding PC model (see table 3) and T_1 model around $\boldsymbol{\mu}_{\mathbf{X}} = \mathbb{E}[\mathbf{X}]$, estimated with a sample of size 1000.

	Y_3	$g_3^{\text{PC}}(\mathbf{X})$	$g_3^{T_1}(\mathbf{X})$
Y_3	1.00	0.80	0.57
$g_3^{\text{PC}}(\mathbf{X})$	0.80	1.00	0.64
$g_3^{T_1}(\mathbf{X})$	0.57	0.64	1.00

Table 5 shows that $g_3^{\text{PC}}(\mathbf{X})$ is well correlated with $Y_3 = f_3(\mathbf{X})$, with a Pearson coefficient of 0.8, while $g_3^{T_1}$ is less so, with a coefficient of 0.57, which may be explained by the strong non-linearity of f_3 . According to eq. (2.7), these correlation coefficients lead to a theoretical variance reduction factor R^2 of 64% when using $g_3^{\text{PC}}(\mathbf{X})$ as a single CV, and of about 32% when using $g_3^{T_1}(\mathbf{X})$ as a single CV. Using both CVs results in a reduction factor of 65%. This minor increase in R^2 can be explained by the modest correlation coefficient of 0.64 between $g_3^{T_1}(\mathbf{X})$ and $g_3^{\text{PC}}(\mathbf{X})$. Note that a reduction factor of R^2 in the variance corresponds to a reduction factor of $1 - \sqrt{1 - R^2}$ in the standard deviation.

These theoretical expectations are reflected in fig. 1a with an RMSE reduction of about 20% when using $g_3^{T_1}$ and 40% when using g_3^{PC} alone or jointly with $g_3^{T_1}$. This figure confirms that g_3^{PC} provides a better CV than $g_3^{T_1}$, reducing the RMSE of the MC estimator twice as much, regardless of the computational budget. However, the construction cost of the surrogate is not the same. While constructing $g_3^{T_1}$ requires only one evaluation of f_3 and its Jacobian matrix, namely at $\boldsymbol{\mu}_{\mathbf{X}}$, the construction of g_3^{PC} involved 100 f_3 evaluations. In the case where the surrogate is built specifically for the estimation of the statistic, the real estimation cost is higher as it includes this construction cost. Figure 1b illustrates this difference by including the surrogate construction cost in the total evaluation cost. As a result, a significant offset appears when using g_3^{PC} , since part of the computational budget (namely 100) is used for the surrogate construction and thus not for the estimation. Specifically, the RMSEs of the CV estimators using g_3^{PC} get below that of the CV estimator using only $g_3^{T_1}$ for a budget of 300 f_3 evaluations. For budgets under 200, using only $g_3^{T_1}$ is preferable, even without the analytical gradient available, which would then incur a construction cost of 8 f_3 evaluations to approximate the gradient using finite differences. This figure also illustrates proposition 2.1, that is, increasing the number of control variates improves (rigorously speaking, does not deteriorate) the variance of the CV estimator.



(a) The evaluation cost does not include the cost of constructing the surrogate models. (b) The evaluation cost includes the cost of constructing the surrogate models.

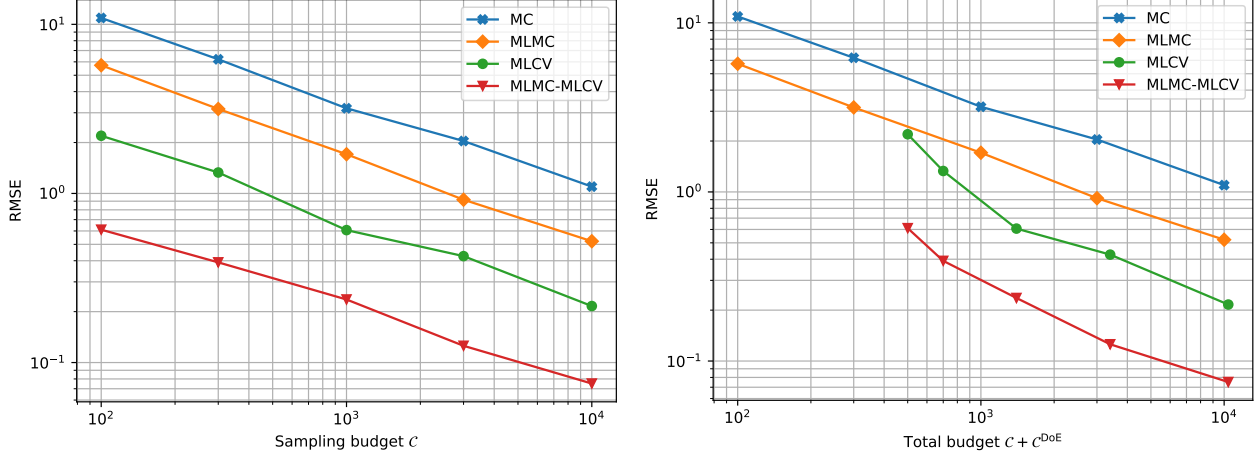
Figure 1: RMSE of the MC and CV estimators of $\theta = \mathbb{E}[f_3(\mathbf{X})]$ (see eq. (4.14)) with respect to the number $\mathcal{C} \in \{100; 300; 1000; 3000; 10\,000\}$ of f_3 evaluations. The CV uses either a first-order Taylor approximation (T1) of f_3 or a PC approximation of f_3 trained from 100 evaluations ($\mathcal{C}^{\text{DoE}} = 100$), or both. The RMSEs are computed using 500 replicates.

4.3.2 MLCV and MLMC-MLCV

Figure 2 compares the MLCV and MLMC-MLCV estimators proposed in eqs. (3.9) and (3.20) with the classical MC and MLMC estimators. This comparison is repeated for different budgets \mathcal{C} , expressed in terms of the equivalent number of f_3 evaluations. Note that the MC and MLCV estimators only use evaluations of the finest simulator f_3 , so that $\mathcal{C} = n_3$, while MLMC-based estimators use all the simulators, so that \mathcal{C} is given in terms of the MLMC cost eq. (3.6), indeed corresponding to the equivalent number of f_3 evaluations, since the costs are normalized such that $\mathcal{C}_3 = 1$. From here on, only PC-based surrogates will be used, so that we omit the superscript “PC” in the notations of the surrogate models.

MLMC vs. MC Figure 2a shows that, for a given sampling budget \mathcal{C} , the MC estimator is the least accurate. This can be explained by the fact that it only has access to the finest simulator, f_3 , whose cost only allows a limited number of evaluations. On the contrary, the MLMC estimator spreads this sampling budget over the four simulators. Ideally, the optimal sample allocation of MLMC eq. (3.35) results in many coarse, cheap evaluations and few fine, expensive evaluations. This is typically the case when the outputs of the simulators are highly correlated and the associated computational cost grows exponentially. Table 7 shows that the first assumption holds. The second assumption also holds since $\mathcal{C}_\ell = \mathcal{O}(KN_\ell) = \mathcal{O}(N_\ell)$, K being fixed, i.e., the evaluation cost grows linearly with the number of quadrature nodes. As a consequence, the MLMC estimator has a lower RMSE than the standard MC estimator.

MLCV vs. MLMC For this experiment, the MLCV estimator is more accurate than the MLMC estimator. Based on one control variate per level ℓ , based on the PC model g_ℓ of f_ℓ from table 3, this estimator dedicates all the sampling budget to the finest simulator f_3 , and uses these control variates based on g_0, g_1, g_2 and g_3 to reduce the variance of the MC estimator at no extra cost, as the evaluation cost of a PC model g_ℓ is negligible compared to the evaluation cost of f_3 . This MLCV



(a) The evaluation cost does not include the cost of construction of the surrogate models. (b) The evaluation cost includes the cost of construction of the surrogate models.

Figure 2: RMSE of the MC, MLMC, MLCV and MLMC-MLCV estimators of $\theta = \mathbb{E}[f_3(\mathbf{X})]$ (see eq. (4.14)) with respect to the sampling budget $c \in \{100; 300; 1000; 3000; 10\,000\}$ of f_3 evaluations. The surrogates used for the MLCV and MLMC-MLCV estimators are described in tables 3 and 4, respectively, and their total construction cost is $c^{\text{DoE}} = 400$. The RMSEs are estimated using 500 replicates.

Table 6: Pearson correlation coefficients between the finest output $Y_3 = f_3(\mathbf{X})$ and the corresponding control variates based on the PC models of table 3 used for the MLCV method, estimated with a sample of size 1000.

	Y_3	$g_0(\mathbf{X})$	$g_1(\mathbf{X})$	$g_2(\mathbf{X})$	$g_3(\mathbf{X})$
Y_3	1.00	0.96	0.97	0.93	0.80
$g_0(\mathbf{X})$	0.96	1.00	0.96	0.92	0.80
$g_1(\mathbf{X})$	0.97	0.96	1.00	0.95	0.81
$g_2(\mathbf{X})$	0.93	0.92	0.95	1.00	0.80
$g_3(\mathbf{X})$	0.80	0.96	0.81	0.80	1.00

technique works particularly well in this case because the control variates are highly correlated to f_3 . Indeed, table 6 shows that their Pearson coefficients are at least 0.8, which guarantees a theoretical reduction of at least 94% in the variance of the MC estimator (i.e. a reduction of at least about 76% in standard deviation), corresponding to the variance reduction when using a single control variate based on g_3 . In fact, the variance reduction factor R^2 when using all the surrogates is only slightly higher, namely $R^2 \approx 95\%$ corresponding to a standard reduction factor of around 78%, which is reflected in fig. 2a.

MLMC-MLCV vs. MLCV Combining the MLMC and MLCV techniques allows the resulting MLMC-MLCV estimator to reduce the variance even more significantly. This can be explained by the very high correlation between Y_0, Y_1, Y_2 and Y_3 on the one hand, which ensures the good performance of the MLMC approach, and by the strong correlation between the control variates on the other hand, ensuring their good performance in combination with the MLMC technique. In

Table 7: Pearson correlation coefficients between the outputs $Y_\ell = f_\ell(\mathbf{X})$ of the simulators and $g_\ell(\mathbf{X})$ of the PC models, as well as the successive differences $Y_\ell - Y_{\ell-1}$ and the outputs $h_\ell(\mathbf{X})$ of the associated PC models, for the MLMC-based estimators. The coefficients are estimated from a sample of size 1000.

	Y_0	Y_1	Y_2	Y_3	$Y_1 - Y_0$	$Y_2 - Y_1$	$Y_3 - Y_2$	$g_0(\mathbf{X})$	$g_1(\mathbf{X})$	$g_2(\mathbf{X})$	$g_3(\mathbf{X})$	$h_1(\mathbf{X})$	$h_2(\mathbf{X})$	$h_3(\mathbf{X})$
Y_0	1.00	0.97	0.97	0.97	-0.11	-0.04	-0.04	0.99	0.96	0.93	0.47	-0.12	-0.13	-0.10
Y_1	0.97	1.00	1.00	1.00	0.13	0.19	0.19	0.97	0.97	0.93	0.48	0.12	0.09	0.09
Y_2	0.97	1.00	1.00	1.00	0.14	0.20	0.20	0.96	0.97	0.92	0.48	0.13	0.10	0.10
Y_3	0.97	1.00	1.00	1.00	0.14	0.20	0.20	0.96	0.97	0.92	0.48	0.13	0.10	0.10
$Y_1 - Y_0$	-0.11	0.13	0.14	0.14	1.00	0.96	0.96	-0.10	0.07	0.01	0.03	0.99	0.90	0.80
$Y_2 - Y_1$	-0.04	0.19	0.20	0.20	0.96	1.00	1.00	-0.03	0.12	0.06	0.06	0.94	0.90	0.81
$Y_3 - Y_2$	-0.04	0.19	0.20	0.20	0.96	1.00	1.00	-0.03	0.12	0.06	0.06	0.95	0.90	0.81
$g_0(\mathbf{X})$	0.99	0.97	0.96	0.96	-0.10	-0.03	-0.03	1.00	0.96	0.93	0.48	-0.11	-0.12	-0.09
$g_1(\mathbf{X})$	0.96	0.97	0.97	0.97	0.07	0.12	0.12	0.96	1.00	0.94	0.47	0.06	0.04	0.04
$g_2(\mathbf{X})$	0.93	0.93	0.92	0.92	0.01	0.06	0.06	0.93	0.94	1.00	0.47	-0.00	-0.03	-0.02
$g_3(\mathbf{X})$	0.47	0.48	0.48	0.48	0.03	0.06	0.06	0.48	0.47	0.47	1.00	0.03	0.01	0.02
$h_1(\mathbf{X})$	-0.12	0.12	0.13	0.13	0.99	0.94	0.95	-0.11	0.06	-0.00	0.03	1.00	0.91	0.80
$h_2(\mathbf{X})$	-0.13	0.09	0.10	0.10	0.90	0.90	0.90	-0.12	0.04	-0.03	0.01	0.91	1.00	0.80
$h_3(\mathbf{X})$	-0.10	0.09	0.10	0.10	0.80	0.81	0.81	-0.09	0.04	-0.02	0.02	0.80	0.80	1.00

particular, table 7 shows that $g_0(\mathbf{X})$, $g_1(\mathbf{X})$ and $g_2(\mathbf{X})$ are highly correlated with Y_3 , with Pearson correlation coefficients greater than 0.9, while $g_3(\mathbf{X})$ is poorly correlated with Y_3 , with a correlation coefficient of 0.48. Besides, $h_1(\mathbf{X})$, $h_2(\mathbf{X})$ and $h_3(\mathbf{X})$ are well-correlated with $Y_1 - Y_0$, $Y_2 - Y_1$ and $Y_3 - Y_2$, respectively, with correlation coefficients greater than 0.8.

A priori estimation of the variance reduction factor and sample allocation Further insights regarding the expected variance reduction of the MLMC-* estimators can be drawn from table 8, which reports the variance reduction factor R_ℓ^2 w.r.t. pure MLMC on each level defined in eqs. (3.25) and (3.26), as well as the quantity \mathcal{S}_ℓ^2 defined in eq. (3.35). In addition, the variance of the MLMC-* estimators of the expectation with optimal sample allocation is S_L^2/\mathcal{C} (see eq. (3.36)), so that the variance reduction factor of an MLMC-* estimator compared to the high-fidelity MC estimator is $1 - \mathcal{S}_L^2/(\mathcal{C}_L \mathbb{V}[Y_L])$, which is reported in the last rows and last column of the table. Consequently, the variance of the different MLMC-based estimators per unit cost can be compared directly through the ratio of their respective S_L^2 . For instance, we observe that the variance of the MLMC-MLCV estimator is significantly reduced compared to that of the MLMC estimator, by about $1 - 40.04/2797.65 \approx 98.6\%$, resulting in a reduction in standard deviation of about 88%. Again, this is well reflected in fig. 2a. Table 8 also reports the share of computational work across correction levels, namely $n_\ell^*(\mathcal{C}_\ell + \mathcal{C}_{\ell-1})/\mathcal{C}$, with n_ℓ^* as defined in eq. (3.35), for the correction on level ℓ . These anticipated shares are consistent with those obtained by using algorithm 1, reported in fig. 4b. It is important to note that the bottom rows of table 8 are deduced from estimates of $\mathbb{V}[Y_\ell]$, \mathcal{V}_ℓ and R_ℓ^2 , reported in the top rows. This means that, in practice, these estimates can be obtained from a few evaluations of f_0, \dots, f_L (and of the appropriate surrogate models), to anticipate the expected variance reduction factor of the MLMC-* estimators and apply algorithm 1 to the most promising.

Remark 4.1. Table 8 highlights the fast decay of \mathcal{V}_ℓ with ℓ , which is a favorable scenario for MLMC. In fact, table 8 reports that the variance of the plain MLMC estimator is reduced by about 75% compared to that of the crude, high-fidelity MC estimator, i.e., a reduction by a factor 4.

Equivalently, this means that a 4-fold increase of the computational budget is needed to achieve the same variance with the high-fidelity MC estimator as that of the plain MLMC. This favorable setting for MLMC makes it all the more challenging to further reduce the variance. Yet, we observe that the proposed MLMC-* estimators are able to achieve a much smaller variance, especially by reducing the MLMC variance on the coarsest levels.

Table 8: Relevant quantities for the MLMC-based estimators, with n_ℓ^* and \mathcal{S}_ℓ^2 as defined in eq. (3.35), \mathcal{C}_ℓ as reported in table 2, and where \mathcal{V}_ℓ and R_ℓ^2 have been estimated using an independent sample of size 10 000. The bold percentages reported in the last rows and last column correspond to the variance reduction factor of the MLMC-* estimators compared to the high-fidelity MC estimator.


ℓ	0	1	2	3
$\mathbb{V}[Y_\ell]$	1.0850×10^4	1.1153×10^4	1.1205×10^4	1.1218×10^4
$\mathcal{V}_\ell = \mathbb{V}[Y_\ell - Y_{\ell-1}]$	1.0850×10^4	5.9029×10^2	1.0590	5.8160×10^{-2}
R_ℓ^2				
MLMC	0	0	0	0
MLMC-CV	0.9838	0.9916	0.8224	0.6469
MLMC-MLCV	0.9840	0.9920	0.9173	0.9202
MLMC-CV[0]	0.9838	0	0	0
MLMC-MLCV[0]	0.9840	0.9916	0.8992	0.9027
$n_\ell^* \frac{\mathcal{C}_\ell + \mathcal{C}_{\ell-1}}{\mathcal{C}}$				
MLMC	69.63 %	28.13 %	1.68 %	0.56 %
MLMC-CV	71.00 %	20.66 %	5.69 %	2.66 %
MLMC-MLCV	73.66 %	20.97 %	4.05 %	1.32 %
MLMC-CV[0]	22.59 %	71.69 %	4.29 %	1.42 %
MLMC-MLCV[0]	72.84 %	21.30 %	4.42 %	1.44 %
\mathcal{S}_ℓ^2				
MLMC	1356.31	2673.54	2766.49	2797.65
MLMC-CV	21.99	36.64	41.33	43.62
MLMC-MLCV	21.72	35.85	38.99	40.04
MLMC-CV[0]	21.99	382.87	418.54	430.71
MLMC-MLCV[0]	21.76	36.35	39.84	41.01
$1 - \frac{\mathcal{S}_\ell^2}{\mathcal{C}_\ell \mathbb{V}[Y_\ell]}$				
MLMC	0.00 %	4.11 %	50.62 %	75.06 %
MLMC-CV	98.38 %	98.69 %	99.26 %	99.61 %
MLMC-MLCV	98.40 %	98.71 %	99.30 %	99.64 %
MLMC-CV[0]	98.38 %	86.27 %	92.53 %	96.16 %
MLMC-MLCV[0]	98.40 %	98.70 %	99.29 %	99.63 %

Accounting for the construction cost of the surrogate model These first results from fig. 2a highlight the interest of multilevel control variates, be it with the MLMC-MLCV estimator or simply with the MLCV one. These results suppose that the PC models are not built specifically for the study, so that the budget does not include the number of f_3 -equivalent simulations required for their construction. Figure 2b illustrates the alternative case where the cost of the surrogate based CV estimators includes the cost of constructing the surrogates. As a result, an additional budget of $\mathcal{C}^{\text{DoE}} = 400$ (see section 4.2), is allocated to the construction of the PC models. As was the case for the single-level PC-based CV estimators, this results in an offset of 400 in the total cost of the MLCV and MLMC-MLCV estimators. The effect is especially noticeable when the construction

budget larger than the estimation budget, i.e. for $\mathcal{C} \in \{100; 300\}$. For a sampling cost of 100, i.e. a total evaluation cost of 500, the MLCV estimator is still slightly more accurate than the MLMC estimator, while the MLMC-MLCV estimator still largely outperforms both.

4.3.3 Variants of MLMC-based CV estimators

The discussion about the surrogate construction budget prompts us to investigate variants of MLMC-MLCV using fewer surrogates, in order to reduce the total evaluation cost for limited budgets. The MLMC-CV estimator is described in section 3.2.2, while the variants MLMC-CV[0] and MLMC-MLCV[0] of MLMC-CV and MLMC-MLCV are introduced in table 1. The MLMC-CV[0] estimator only uses a control variate based on g_0 , so that the construction cost drops to $\mathcal{C}^{\text{DoE}} = 100$, while the MLMC-MLCV[0] estimator uses control variates based on g_0 , g_1 and h_1 , so that the construction cost drops to $\mathcal{C}^{\text{DoE}} = 200$. The MLMC-CV estimator uses control variates based on surrogates at all levels, so that the construction cost remains $\mathcal{C}^{\text{DoE}} = 400$.



(a) The evaluation cost does not include the cost of construction of the surrogate models.

(b) The evaluation cost includes the cost of construction of the surrogate models.

Figure 3: RMSE of the MLMC-MLCV, MLMC-CV, MLMC-MLCV[0], and MLMC-CV[0] estimators of $\theta = \mathbb{E}[f_3(\mathbf{X})]$ (see eq. (4.14)) with respect to the sampling budget $\mathcal{C} \in \{100; 300; 1000; 3000; 10\,000\}$ of f_3 evaluations. The surrogates used for the MLMC-based estimators are described in table 4. The RMSEs of the pure PC estimators, summarized in table 9 are shown in black for comparison. The RMSEs for the MLMC-* estimators are computed using 500 replicates, while the RMSEs of the PC-only estimators are computed from 200 replicates.

Figure 3a shows that MLMC-CV[0] has much higher RMSE than the other variants, resulting from the fact that it only reduces the variance associated with the coarsest level of the MLMC estimator. This behavior is consistent with the quantities of table 8. In particular, the value of S_L^2 is about 10 times higher than for the other variants, accounting for its RMSE being about 3 times higher than for the other variants. The remaining variants have similar performances regardless of the estimation budget, which is consistent with the S_L^2 values given in table 8. On the other hand, when considering the construction cost of the surrogates, fig. 3b shows that MLMC-MLCV[0] performs best, as it uses only surrogate models related to the two coarsest levels, namely g_0 , g_1 and h_1 , so that the construction cost is reduced. Furthermore, these surrogates have excellent Q^2 , and they are such that $g_0(\mathbf{X})$ is highly correlated with Y_0 , $g_1(\mathbf{X})$ is highly correlated with Y_1 , and $h_1(\mathbf{X})$ is highly correlated with $Y_1 - Y_0$, $Y_2 - Y_1$ and $Y_3 - Y_2$. Namely, the associated Pearson

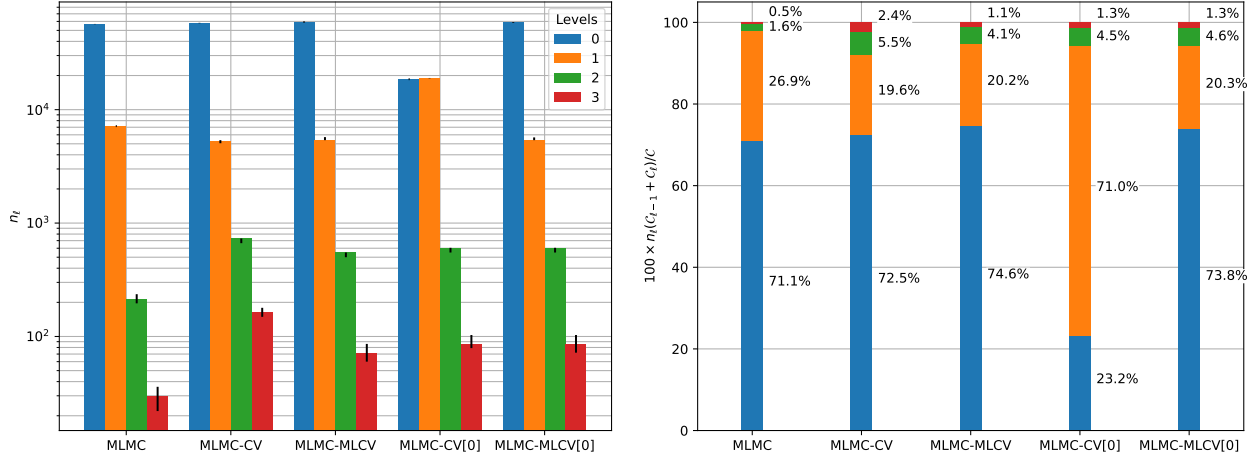
correlation coefficients reported in table 7 are all at least 0.94. Therefore, should one have to build the surrogates specifically for the CV estimation of a statistic, it is more advantageous to adopt the MLMC-MLCV[0] variant over the others. In our case, the construction budget is divided by two compared to MLMC-MLCV and MLMC-CV, for a similar performance in terms of RMSE. For completeness, we also compare the proposed estimators with plain PC models of the high-fidelity simulator f_3 , whose characteristics are summarized in table 9. The black curve in fig. 3 shows that, for the same budget, combining MLMC sampling and multifidelity surrogate models is more advantageous than using a surrogate model of the highest fidelity model alone. By extrapolating, the latter should eventually become better than the MLMC-* estimators around a budget of 3000, which is much higher than practically affordable for expensive simulators. Note also that the PC models used here for comparison are optimized using the entire budget, while the MLMC-* estimators use PC models that are constructed from an arbitrary and non-optimized distribution of the sampling budget between the different model levels as reported in table 4. Strategies for optimizing this tradeoff, which constitutes a promising avenue for further improvement, will be investigated in future work.

Table 9: PC models of the high-fidelity simulator f_L used directly for the estimation of $\mathbb{E}[\mathcal{M}(\mathbf{X})]$, with their sample size, degree and quality measure. These models are built from an optimized LHS DoE, using the basis-adaptive LARS algorithm of [7, Fig. 5]. The average values of p^* , $|\tilde{\mathcal{A}}_{p^*}|$ and Q^2 over 200 replications is reported, along with their respective standard deviation between brackets.

n_L^{DoE}	100	200	500	1000
p^*	4(2)	6(1)	9(1)	11(1)
$ \tilde{\mathcal{A}}_{p^*} $	20(13)	48(16)	118(31)	219(56)
Q^2	0.35(21)	0.84(2)	0.95(1)	0.98

4.3.4 Budget allocation

Lastly, fig. 4a shows the number n_ℓ of evaluations for each of the $L + 1$ correction levels of the MLMC-* telescopic sum. Precisely, n_0 is the number of evaluations of f_0 , while n_ℓ is the number of evaluations of $f_\ell - f_{\ell-1}$, for $\ell > 0$. We observe a typical sample allocation for MLMC-like estimators in ideal cases, that is, many coarse evaluations, and fewer and fewer fine evaluations. The MLMC-CV[0] estimator slightly deviates from this pattern, with $n_1 \approx n_0$, which can be explained by the fact that $\mathcal{V}_1 \approx 590$ is of the same order of magnitude as $(1 - R_0^2)\mathcal{V}_0 \approx 176$. Figure 4b depicts the share of overall sampling cost associated with the different correction levels. Specifically, $n_0\mathcal{C}_0\mathcal{C}^{-1}$ is the share of sampling budget dedicated to evaluating f_0 , and $n_\ell(\mathcal{C}_\ell + \mathcal{C}_{\ell-1})\mathcal{C}^{-1}$ is the share dedicated to evaluating $f_\ell - f_{\ell-1}$, for $\ell > 0$. We see that, except for MLMC-CV[0] for the reasons explained above, most of the sampling budget (around 70%) is allocated to the coarsest level, and most of the remaining budget is dedicated to correction level $\ell = 1$. We note that the CV-based MLMC estimators dedicate slightly more budget to levels 2 and 3 than for the standard MLMC estimator. The sampling budget shares of fig. 4b are consistent with the theoretical optimal shares reported in table 8, including those of the MLMC-CV[0] estimator. This suggests that the sample allocation resulting from algorithm 1 seems to converge to the theoretical optimal sample allocation eq. (3.35).



(a) Sample allocation. The bars represent the median sample size n_ℓ associated with the corresponding correction level ℓ of the MLMC-* telescopic sum. The black vertical lines represent the 25% and 75% quantiles.

(b) Budget allocation. The stacked bars represent the proportion of the sampling cost dedicated to each level of the MLMC-* telescopic sum.

Figure 4: Sample allocation and associated computational cost across the correction levels for the MLMC-based estimators of $\theta = \mathbb{E}[f_3(\mathbf{X})]$ (see eq. (4.14)), with an estimation budget of $\mathcal{C} = 10\,000$.

5 Conclusions

In this paper, we proposed multilevel variance reduction strategies relying on surrogate-based control variates. On the one hand, using specific surrogate models, such as polynomial chaos expansions or Taylor polynomial expansions, allows to directly access exact statistics (mean, variance) of the control variates. Even if these exact statistics are not directly accessible (e.g., when using GPs), they can be estimated very accurately at negligible cost. This contrasts with typical control variates relying on lower-fidelity models based on models/simulators with degraded physics or coarser discretizations, for which approximate control variate strategies need to be devised, resulting in a lower variance reduction. On the other hand, when multiple levels of fidelities (e.g., based on the discretization) with a clear cost/accuracy hierarchy are available, the surrogate-based control variate approach can be efficiently combined with multilevel strategies. The first strategy, MLCV, simply consists of using multiple control variates based on surrogate models of the simulators corresponding to the different levels. The main advantage is that the surrogate models corresponding to coarse levels may be constructed using larger sample sizes than for the finest level, resulting in a more accurate surrogate model (i.e., with lower model error). This strategy thus leads to a greater variance reduction compared to only using one surrogate model based on the finest level. This is supported by the numerical experiments we conducted, as well as by the theoretical variance reduction provided by proposition 2.1. The second strategy, MLMC-MLCV, allows to further improve the variance reduction by combining the surrogate-based control variates with an MLMC strategy. The most appropriate way to construct and utilize the surrogate models is, however, not straightforward, and was discussed in detail in sections 3.2.2 and 3.2.3. The additional variance reduction as compared to plain MLMC is demonstrated in our numerical experiments and supported by the theoretical variance reduction factor eq. (3.26) derived in section 3.2.3.

The construction cost of the surrogate models was discussed from two perspectives. When the surrogate models are constructed for the sole purpose of serving for the control variate estimation,

then the cost of their construction must be taken into account for fair comparison with other approaches. In such a case, it may not be optimal, in terms of cost/accuracy tradeoff, to construct surrogate models on all levels, especially when only a limited budget is available. In particular, using a subset of surrogate models based on the coarser levels may already lead to considerable variance reduction, provided that the outputs of the coarse surrogate models are sufficiently correlated with that of the high-fidelity simulator. Then, considering additional surrogate models on finer levels might only result in marginal improvement, at the expense of a significant computational cost. On the contrary, if the surrogate models have already been constructed for other purposes, and are, in some sense, available “for free,” then their construction cost need not be considered, and the entire set of surrogate models may then be used.

From the former perspective, that is when the cost of the surrogate construction is considered as part of the estimation cost, one may devise more involved strategies seeking to optimize the tradeoff between the construction cost and the model error, directly impacting the projected variance reduction. For instance, in the context of polynomial chaos surrogate models, the truncation strategy may be controlled to this end, as proposed in a stochastic Galerkin framework in [56], where the total polynomial degree is optimized alongside the sample size and the CV parameter to minimize the PC-based CV estimator’s variance under a cost constraint. Another avenue to improve the proposed approach would be to replace the MLMC part of the MLMC-MLCV strategy by a more efficient multilevel approach, such as the multilevel best linear unbiased estimator (MLBLUE) [48, 49, 47]. In particular, an MLBLUE-MLCV strategy should be more efficient when a collection of low-fidelity simulators (e.g., with degraded physics) with no clear cost/accuracy hierarchy is available. Finally, although the proposed approaches apply to the estimation of arbitrary statistics, they were only tested here on the estimation of expected values. The theoretical and algorithmic ingredients for the estimation of variances are, however, described in this paper and may be tested in follow-up investigations and numerical experiments. Specifically, the multifidelity estimation of variance-based sensitivity indices is of particular interest to our team.

6 Acknowledgment

We wish to acknowledge the PIA framework (CGI, ANR) and the industrial members of the IRT Saint Exupéry project R-Evol: Airbus, Liebherr, Altran Technologies, Capgemini DEMS France, CENAERO and Cerfacs for their support, financial funding and own knowledge.

A Optimal parameter for expectation and variance CV estimators

We derive here the expression of the optimal CV parameter $\boldsymbol{\alpha}^*$ for the CV estimators of the expectation and of the variance. We define $Y = f(\mathbf{X})$, $Z_m = g_m(\mathbf{X})$ and $\mathbf{Z} = (Z_m)_{m=1}^M$. Then, given an input n -sample $\{\mathbf{X}^{(i)}\}_{i=1}^n$, we define $Y^{(i)} = f(\mathbf{X}^{(i)})$, $Z_m^{(i)} = g_m(\mathbf{X}^{(i)})$ and $\mathbf{Z}^{(i)} = (Z_m^{(i)})_{m=1}^M$.

A.1 CV estimator of the expectation

For the expectation, we have $\boldsymbol{\Sigma} = \mathbb{C}[\hat{E}[\mathbf{Z}]]$ and $\mathbf{c} = \mathbb{C}[\hat{E}[Y], \hat{E}[\mathbf{Z}]]$, i.e.

$$[\boldsymbol{\Sigma}]_{m,m'} = \mathbb{C}[\hat{E}[Z_m], \hat{E}[Z_{m'}]] = n^{-2} \sum_{i,j=1}^n \mathbb{C}[Z_m^{(i)}, Z_{m'}^{(j)}] = n^{-1} \mathbb{C}[Z_m, Z_{m'}], \quad (\text{A.1})$$

$$[\mathbf{c}]_m = \mathbb{C}[\hat{E}[Y], \hat{E}[Z_m]] = n^{-2} \sum_{i,j=1}^n \mathbb{C}[Y^{(i)}, Z_m^{(j)}] = n^{-1} \mathbb{C}[Y, Z_m], \quad (\text{A.2})$$

so that $\boldsymbol{\Sigma} = n^{-1} \mathbb{C}[\mathbf{Z}]$, $\mathbf{c} = n^{-1} \mathbb{C}[Y, \mathbf{Z}]$, and, eventually,

$$\boldsymbol{\alpha}^* = \mathbb{C}[\mathbf{Z}]^{-1} \mathbb{C}[Y, \mathbf{Z}]. \quad (\text{A.3})$$

Furthermore, we have $\mathbb{V}[\hat{E}[Y]] = n^{-1} \mathbb{V}[Y]$, so that

$$R^2 = \frac{\mathbb{C}[Y, \mathbf{Z}]^\top \mathbb{C}[\mathbf{Z}]^{-1} \mathbb{C}[Y, \mathbf{Z}]}{\mathbb{V}[Y]} = \mathbf{r}_{Y,\mathbf{Z}}^\top \mathbf{R}_{\mathbf{Z}}^{-1} \mathbf{r}_{Y,\mathbf{Z}}, \quad (\text{A.4})$$

where

$$\mathbf{r}_{Y,\mathbf{Z}} = (\mathbb{V}[Y] \mathbf{D}_{\mathbf{Z}})^{-1/2} \mathbb{C}[Y, \mathbf{Z}], \quad \mathbf{R}_{\mathbf{Z}} = \mathbf{D}_{\mathbf{Z}}^{-1/2} \mathbb{C}[\mathbf{Z}] \mathbf{D}_{\mathbf{Z}}^{-1/2}, \quad \mathbf{D}_{\mathbf{Z}} = \text{Diag}(\mathbb{C}[\mathbf{Z}]). \quad (\text{A.5})$$

Thus, $R^2 \in [0, 1]$ corresponds to the squared coefficient of multiple correlation between Y and the control variates Z_1, \dots, Z_M .

A.2 CV estimator of the variance

Similarly, for the variance, we have $\boldsymbol{\Sigma} = \mathbb{C}[\hat{V}[\mathbf{Z}]]$ and $\mathbf{c} = \mathbb{C}[\hat{V}[Y], \hat{V}[\mathbf{Z}]]$. We start by deriving useful identities. In what follows, for any random variable A , we denote the corresponding centered variable by $\bar{A} := A - \mathbb{E}[A]$. First, we remark that $\hat{V}[A] = \hat{V}[\bar{A}]$, so that, for any two random variables Y and Z ,

$$\mathbb{C}[\hat{V}[Y], \hat{V}[Z]] = \mathbb{E}[\hat{V}[\bar{Y}] \hat{V}[\bar{Z}]] - \mathbb{V}[Y] \mathbb{V}[Z]. \quad (\text{A.6})$$

Furthermore, it can be shown that

$$\mathbb{E}[\hat{V}[\bar{Y}] \hat{V}[\bar{Z}]] = \left(\frac{n}{n-1} \right)^2 (a_n(\bar{Y}, \bar{Z}) + b_n(\bar{Y}, \bar{Z}) - c_n(\bar{Y}, \bar{Z}) - c_n(\bar{Z}, \bar{Y})), \quad (\text{A.7})$$

with (see proof below)

$$a_n(\bar{Y}, \bar{Z}) := \mathbb{E}[\hat{E}[\bar{Y}^2] \hat{E}[\bar{Z}^2]] = \frac{1}{n} \mathbb{C}[\bar{Y}^2, \bar{Z}^2] + \mathbb{V}[Y] \mathbb{V}[Z] = a_n(\bar{Z}, \bar{Y}), \quad (\text{A.8})$$

$$b_n(\bar{Y}, \bar{Z}) := \mathbb{E}[\hat{E}[\bar{Y}^2] \hat{E}[\bar{Z}^2]] = \frac{a_n(\bar{Y}, \bar{Z})}{n^2} + 2 \frac{n-1}{n^3} \mathbb{C}[Y, Z]^2 = b_n(\bar{Z}, \bar{Y}), \quad (\text{A.9})$$

$$c_n(\bar{Y}, \bar{Z}) := \mathbb{E}[\hat{E}[\bar{Y}^2] \hat{E}[\bar{Z}^2]] = \frac{a_n(\bar{Y}, \bar{Z})}{n} = c_n(\bar{Z}, \bar{Y}). \quad (\text{A.10})$$

We thus have

$$\mathbb{C}[\hat{V}[Y], \hat{V}[Z]] = \frac{1}{n}\mathbb{C}[\bar{Y}^2, \bar{Z}^2] + \frac{2}{n(n-1)}\mathbb{C}[Y, Z]^2, \quad (\text{A.11})$$

$$\mathbb{V}[\hat{V}[Y]] = \mathbb{C}[\hat{V}[Y], \hat{V}[Y]] = \frac{1}{n}\mathbb{V}[\bar{Y}^2] + \frac{2}{n(n-1)}\mathbb{V}[Y]^2, \quad (\text{A.12})$$

eventually leading to

$$\boldsymbol{\Sigma} = \mathbb{C}[\hat{V}[\mathbf{Z}]] = \frac{1}{n} \left(\mathbb{C}[\bar{\mathbf{Z}}^{\circ 2}] + \frac{2}{n-1}\mathbb{C}[\mathbf{Z}]^{\circ 2} \right), \quad (\text{A.13})$$

$$\mathbf{c} = \mathbb{C}[\hat{V}[Y], \hat{V}[\mathbf{Z}]] = \frac{1}{n} \left(\mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\circ 2}] + \frac{2}{n-1}\mathbb{C}[Y, \mathbf{Z}]^{\circ 2} \right), \quad (\text{A.14})$$

so that

$$\boldsymbol{\alpha}^* = \left[\mathbb{C}[\bar{\mathbf{Z}}^{\circ 2}] + \frac{2}{n-1}\mathbb{C}[\mathbf{Z}]^{\circ 2} \right]^{-1} \left[\mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\circ 2}] + \frac{2}{n-1}\mathbb{C}[Y, \mathbf{Z}]^{\circ 2} \right], \quad (\text{A.15})$$

$$R^2 = \frac{\left[\mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\circ 2}] + \frac{2}{n-1}\mathbb{C}[Y, \mathbf{Z}]^{\circ 2} \right]^{\top} \boldsymbol{\alpha}^*}{\mathbb{V}[\bar{Y}^2] + \frac{2}{n-1}\mathbb{V}[Y]^2}. \quad (\text{A.16})$$

As $n \rightarrow \infty$, we see that

$$\boldsymbol{\alpha}^* \rightarrow \mathbb{C}[\bar{\mathbf{Z}}^{\circ 2}]^{-1}\mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\circ 2}] \quad (\text{A.17})$$

$$R^2 \rightarrow \mathbb{V}[\bar{Y}^2]^{-1}\mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\circ 2}]^{\top}\mathbb{C}[\bar{\mathbf{Z}}^{\circ 2}]^{-1}\mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\circ 2}] = \mathbf{r}_{Y, \mathbf{Z}}^{\top} \mathbf{R}_{\mathbf{Z}}^{-1} \mathbf{r}_{Y, \mathbf{Z}} =: R_{\text{lim}}^2, \quad (\text{A.18})$$

where

$$\mathbf{r}_{Y, \mathbf{Z}} = (\mathbb{V}[\bar{Y}^2]\mathbf{D}_{\mathbf{Z}})^{-1/2}\mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\circ 2}], \quad \mathbf{R}_{\mathbf{Z}} = \mathbf{D}_{\mathbf{Z}}^{-1/2}\mathbb{C}[\bar{\mathbf{Z}}^{\circ 2}]\mathbf{D}_{\mathbf{Z}}^{-1/2}, \quad (\text{A.19})$$

with $\mathbf{D}_{\mathbf{Z}} = \text{Diag}(\mathbb{C}[\bar{\mathbf{Z}}^{\circ 2}])$. Thus, $R_{\text{lim}}^2 \in [0, 1]$ corresponds to the squared coefficient of multiple correlation between \bar{Y}^2 and $\bar{Z}_1^2, \dots, \bar{Z}_M^2$.

We now proceed to the proof of identities eqs. (A.8) to (A.10). First, for eq. (A.8), by definition

$$a_n(\bar{Y}, \bar{Z}) = \frac{1}{n^2}\mathbb{E}[\sum_{i=1}^n (\bar{Y}^{(i)})^2 \sum_{i=1}^n (\bar{Z}^{(i)})^2] = \frac{1}{n^2} \sum_{i,j=1}^n \mathbb{E}[(\bar{Y}^{(i)})^2 (\bar{Z}^{(j)})^2]. \quad (\text{A.20})$$

We distinguish two (disjoint) cases:

1. $i = j$: $\mathbb{E}[(\bar{Y}^{(i)})^2 (\bar{Z}^{(j)})^2] = \mathbb{E}[(\bar{Y}^{(i)})^2 (\bar{Z}^{(i)})^2] = \mathbb{E}[\bar{Y}^2 \bar{Z}^2] = \mathbb{C}[\bar{Y}^2, \bar{Z}^2] + \mathbb{V}[Y]\mathbb{V}[Z]$. There are n such terms in the sum.
2. $i \neq j$: $\mathbb{E}[(\bar{Y}^{(i)})^2 (\bar{Z}^{(j)})^2] = \mathbb{E}[(\bar{Y}^{(i)})^2] \mathbb{E}[(\bar{Z}^{(j)})^2] = \mathbb{E}[\bar{Y}^2] \mathbb{E}[\bar{Z}^2] = \mathbb{V}[Y]\mathbb{V}[Z]$. There are $n(n-1)$ such terms in the sum.

Then eq. (A.8) follows. For eq. (A.9),

$$b_n(\bar{Y}, \bar{Z}) = \mathbb{E}[(\sum_{i=1}^n \bar{Y}^{(i)})^2 (\sum_{i=1}^n \bar{Z}^{(i)})^2] = \frac{1}{n^4} \sum_{i,j,k,\ell=1}^n \mathbb{E}[\bar{Y}^{(i)} \bar{Y}^{(j)} \bar{Z}^{(k)} \bar{Z}^{(\ell)}]. \quad (\text{A.21})$$

We distinguish five (disjoint) cases:

1. $i = j = k = \ell$: $\mathbb{E}[\bar{Y}^{(i)}\bar{Y}^{(j)}\bar{Z}^{(k)}\bar{Z}^{(\ell)}] = \mathbb{E}[(\bar{Y}^{(i)})^2(\bar{Z}^{(i)})^2] = \mathbb{E}[\bar{Y}^2\bar{Z}^2] = \mathbb{C}[\bar{Y}^2, \bar{Z}^2] + \mathbb{V}[Y]\mathbb{V}[Z]$. There are n such terms in the sum.
2. $i = j \neq k = \ell$: $\mathbb{E}[\bar{Y}^{(i)}\bar{Y}^{(j)}\bar{Z}^{(k)}\bar{Z}^{(\ell)}] = \mathbb{E}[(\bar{Y}^{(i)})^2]\mathbb{E}[(\bar{Z}^{(k)})^2] = \mathbb{E}[\bar{Y}^2]\mathbb{E}[\bar{Z}^2] = \mathbb{V}[Y]\mathbb{V}[Z]$. There are $n(n-1)$ such terms in the sum.
3. $i = k \neq j = \ell$: $\mathbb{E}[\bar{Y}^{(i)}\bar{Y}^{(j)}\bar{Z}^{(k)}\bar{Z}^{(\ell)}] = \mathbb{E}[\bar{Y}^{(i)}\bar{Z}^{(i)}]\mathbb{E}[\bar{Y}^{(j)}\bar{Z}^{(j)}] = \mathbb{E}[\bar{Y}\bar{Z}]^2 = \mathbb{C}[Y, Z]^2$. There are $n(n-1)$ such terms in the sum.
4. $i = \ell \neq j = k$: $\mathbb{E}[\bar{Y}^{(i)}\bar{Y}^{(j)}\bar{Z}^{(k)}\bar{Z}^{(\ell)}] = \mathbb{E}[\bar{Y}^{(i)}\bar{Z}^{(i)}]\mathbb{E}[\bar{Y}^{(j)}\bar{Z}^{(j)}] = \mathbb{E}[\bar{Y}\bar{Z}]^2 = \mathbb{C}[Y, Z]^2$. There are $n(n-1)$ such terms in the sum.
5. All remaining cases (at least one of the indices i, j, k, ℓ is different from all the others): $\mathbb{E}[\bar{Y}^{(i)}\bar{Y}^{(j)}\bar{Z}^{(k)}\bar{Z}^{(\ell)}] = 0$.

Then eq. (A.9) follows. Finally, for eq. (A.10),

$$c_n(\bar{Y}, \bar{Z})\mathbb{E}[(\sum_{i=1}^n (\bar{Y}^{(i)})^2)(\sum_{i=1}^n \bar{Z}^{(i)})^2] = \frac{1}{n^3} \sum_{i,j,k=1}^n \mathbb{E}[(\bar{Y}^{(i)})^2 \bar{Z}^{(j)} \bar{Z}^{(k)}]. \quad (\text{A.22})$$

We distinguish three (disjoint) cases:

1. $i = j = k$: $\mathbb{E}[(\bar{Y}^{(i)})^2 \bar{Z}^{(j)} \bar{Z}^{(k)}] = \mathbb{E}[(\bar{Y}^{(i)})^2 (\bar{Z}^{(i)})^2] = \mathbb{E}[\bar{Y}^2 \bar{Z}^2] = \mathbb{M}^4[Y, Z]$. There are n such terms in the sum.
2. $i \neq j = k$: $\mathbb{E}[(\bar{Y}^{(i)})^2 \bar{Z}^{(j)} \bar{Z}^{(k)}] = \mathbb{E}[(\bar{Y}^{(i)})^2]\mathbb{E}[(\bar{Z}^{(j)})^2] = \mathbb{E}[\bar{Y}^2]\mathbb{E}[\bar{Z}^2] = \mathbb{V}[Y]\mathbb{V}[Z]$. There are $n(n-1)$ such terms in the sum.
3. All remaining cases (at least one of the indices j, k is different from all the others): $\mathbb{E}[(\bar{Y}^{(i)})^2 \bar{Z}^{(j)} \bar{Z}^{(k)}] = 0$.

Then eq. (A.10) follows.

Remark A.1. When the expected value $\boldsymbol{\mu}_{\mathbf{Z}}$ of \mathbf{Z} is known, as is the case when defining \mathbf{Z} from the prediction of certain surrogate models, such as PC expansion and Taylor polynomials (and, in some instances, GPs, see section 2.4.1), it is possible to replace $\hat{\mathbf{V}}[\mathbf{Z}]$ with $\hat{\mathbf{E}}[\bar{\mathbf{Z}}^{\odot 2}]$. The derivation of the optimal CV parameter is somewhat easier and leads to similar results as in the unknown expectation case. Specifically,

$$[\boldsymbol{\Sigma}]_{m,m'} = \mathbb{C}[\hat{\mathbf{E}}[\bar{Z}_m^2], \hat{\mathbf{E}}[\bar{Z}_{m'}^2]] = \mathbb{C}[n^{-1} \sum_{i=1}^n (\bar{Z}_m^{(i)})^2, n^{-1} \sum_{i=1}^n (\bar{Z}_{m'}^{(i)})^2] \quad (\text{A.23})$$

$$= n^{-2} \sum_{i,j=1}^n \mathbb{C}[(\bar{Z}_m^{(i)})^2, (\bar{Z}_{m'}^{(j)})^2] = n^{-2} \sum_{i=1}^n \mathbb{C}[(\bar{Z}_m^{(i)})^2, (\bar{Z}_{m'}^{(i)})^2] \quad (\text{A.24})$$

$$= n^{-1} \mathbb{C}[\bar{Z}_m^2, \bar{Z}_{m'}^2], \quad (\text{A.25})$$

i.e. $\boldsymbol{\Sigma} = n^{-1} \mathbb{C}[\bar{\mathbf{Z}}^{\odot 2}]$. Regarding the vector of covariances \mathbf{c} ,

$$[\mathbf{c}]_m = \mathbb{C}[\hat{\mathbf{V}}[\bar{Y}], \hat{\mathbf{E}}[\bar{Z}_m^2]] = \mathbb{E}[\hat{\mathbf{V}}[\bar{Y}]\hat{\mathbf{E}}[\bar{Z}_m^2]] - \mathbb{V}[Y]\mathbb{V}[Z_m] \quad (\text{A.26})$$

$$= \frac{n}{n-1} (\mathbb{E}[\hat{\mathbf{E}}[\bar{Y}^2]\hat{\mathbf{E}}[\bar{Z}_m^2]] - \mathbb{E}[\hat{\mathbf{E}}[\bar{Y}]^2\hat{\mathbf{E}}[\bar{Z}_m^2]]) - \mathbb{V}[Y]\mathbb{V}[Z_m] \quad (\text{A.27})$$

$$= \frac{n}{n-1} (a_n(\bar{Y}, \bar{Z}_m) - c_n(\bar{Z}_m, \bar{Y})) - \mathbb{V}[Y]\mathbb{V}[Z_m] = a_n(\bar{Y}, \bar{Z}_m) - \mathbb{V}[Y]\mathbb{V}[Z_m] \quad (\text{A.28})$$

$$= n^{-1} \mathbb{C}[\bar{Y}^2, \bar{Z}_m^2], \quad (\text{A.29})$$

i.e. $\mathbf{c} = n^{-1}\mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\odot 2}]$, so that

$$\boldsymbol{\alpha}^* = \mathbb{C}[\bar{\mathbf{Z}}^{\odot 2}]^{-1}\mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\odot 2}], \quad (\text{A.30})$$

$$R^2 = \frac{\mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\odot 2}]^\top \mathbb{C}[\bar{\mathbf{Z}}^{\odot 2}]^{-1} \mathbb{C}[\bar{Y}^2, \bar{\mathbf{Z}}^{\odot 2}]}{\mathbb{V}[\bar{Y}^2] + \frac{2}{n-1} \mathbb{V}[Y]^2} \xrightarrow{n \rightarrow \infty} R_{\text{lim}}^2, \quad (\text{A.31})$$

with the same definitions as in eq. (A.18).

B Variance reduction of bi-fidelity (A)CV estimators of the expectation

In this appendix, we illustrate the superiority of the (exact) CV estimator of the expectation over the bi-fidelity ACV estimator introduced in [39, section 3] (see also [19, section IV]), hereafter referred to as the MFMC estimator, in terms of variance reduction. We consider here a high-fidelity simulator f and a low-fidelity version g . We denote by ρ the Pearson correlation coefficient between $f(\mathbf{X})$ and $g(\mathbf{X})$, and we define $w := c_g/c_f$ as the ratio of the expected cost c_g of evaluating $g(\mathbf{X})$ to the expected cost c_f of evaluating $f(\mathbf{X})$. Note that our definition of w is the inverse of the definition of w in [39, 19], allowing our w to remain finite when the expected cost of evaluating $g(\mathbf{X})$ vanishes. We shall assume hereafter that $w, \rho^2 \in (0, 1)$.

B.1 The MFMC estimator

For a given (expected) computational budget \mathcal{C} , the optimal MFMC estimator [39, 19] reads

$$\hat{\mu}_{\mathcal{C}}^{\text{MFMC}} = \frac{1}{N} \sum_{i=1}^N f(\mathbf{X}^{(i)}) - \alpha \left(\frac{1}{N} \sum_{i=1}^N g(\mathbf{X}^{(i)}) - \frac{1}{N'} \sum_{i=1}^{N'} g(\mathbf{X}^{(i)}) \right), \quad (\text{B.1})$$

where

$$\alpha = \frac{\mathbb{C}[Y, Z]}{\mathbb{V}[Z]}, \quad N = \frac{\mathcal{C}}{c_f(1 + \eta w)}, \quad \eta = \sqrt{\frac{\rho^2}{w(1 - \rho^2)}}, \quad N' = \eta N, \quad (\text{B.2})$$

and where $\mathcal{X} := \{\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(N')}\}$ is an N' -sample of \mathbf{X} (we assume here that $N' > N$; see [39] and corollary B.6 for a discussion on this assumption). Note that N and N' need to be appropriately rounded to integers for the estimator eq. (B.1) to make sense. For the same (expected) budget \mathcal{C} , one can afford $p := \mathcal{C}/c_f$ evaluations of $f(\mathbf{X})$, i.e., the expected cost of p evaluations of $f(\mathbf{X})$ is \mathcal{C} . Let $\hat{\mu}_p := p^{-1} \sum_{i=1}^p f(\mathbf{X}^{(i)})$ denote the standard MC estimator of the expected value of $f(\mathbf{X})$ using a p -sample of \mathbf{X} . We now proceed to proving some useful properties of the MFMC estimator eq. (B.1).

Proposition B.1. *The variance of the MFMC estimator defined by eqs. (B.1) and (B.2) is such that $\mathbb{V}[\hat{\mu}_{\mathcal{C}}^{\text{MFMC}}] = \beta^{\text{MFMC}} \mathbb{V}[\hat{\mu}_p]$, with $\beta^{\text{MFMC}} = (1 + \eta w)^2 (1 - \rho^2)$.*

Proof. From [39, Eq. (6)] (see also [19, Eq. (9)]), we have

$$\beta^{\text{MFMC}} = (1 + \eta w) \left[1 - \left(1 - \frac{1}{\eta} \right) \rho^2 \right]. \quad (\text{B.3})$$

Furthermore,

$$1 - \left(1 - \frac{1}{\eta}\right) \rho^2 = 1 - \rho^2 + \frac{\rho^2}{\eta} = (1 - \rho^2) \left(1 + \frac{1}{\eta} \frac{\rho^2}{1 - \rho^2}\right). \quad (\text{B.4})$$

Finally, injecting the expression of η defined in eq. (B.2), we notice that

$$\frac{1}{\eta} \frac{\rho^2}{1 - \rho^2} = \sqrt{\frac{w(1 - \rho^2)}{\rho^2}} \frac{\rho^2}{1 - \rho^2} = \sqrt{\frac{w\rho^2}{1 - \rho^2}} = \eta w, \quad (\text{B.5})$$

which concludes the proof. \square

We note that β^{MFMC} does not depend on \mathcal{C} , and that the smaller β^{MFMC} , the greater the variance reduction with respect to the standard, high-fidelity MC estimator for the same (expected) budget. We also note that variance reduction is achieved if and only if $\beta^{\text{MFMC}} < 1$.

Proposition B.2. $\eta > 1$ (i.e., $N' > N$) if and only if $\rho^2 > w/(1 + w)$.

Proof. $\eta > 1 \iff \eta^2 > 1 \iff \rho^2 > w(1 - \rho^2) \iff (1 + w)\rho^2 > w$, hence the result follows. \square

Proposition B.3. $\beta^{\text{MFMC}} < 1$ if and only if $\rho^2 > 4w/(1 + w)^2$.

Proof. We start by noticing that $1/(1 - \rho^2) = 1 + \rho^2/(1 - \rho^2) = 1 + \eta^2 w$. Then, it follows that $\beta^{\text{MFMC}} < 1 \iff (1 + \eta w)^2 < 1/(1 - \rho^2) = 1 + \eta^2 w \iff [2 - (1 - w)\eta]\eta w < 0$. Because $w, \eta > 0$ and $(1 - w) \in (0, 1)$, we deduce that $\beta^{\text{MFMC}} < 1 \iff \eta > 2/(1 - w) \iff \eta^2 > 4/(1 - w)^2$. Injecting the expression of η defined in eq. (B.2), we get $\beta^{\text{MFMC}} < 1 \iff \rho^2(1 - w)^2 > 4w(1 - \rho^2)$, eventually leading to the desired result. \square

Proposition B.4. For any $w \in (0, 1)$, $4w/(1 + w)^2 < 1$.

Proof. The function $a: x \mapsto 4x/(1 + x)^2$ is continuous on $[0, 1]$, and its derivative is defined on $[0, 1]$ by $a'(x) = 4(1 - x)/(1 + x)^3$, which is positive for $x \in (0, 1)$, so a is strictly increasing on $(0, 1)$. Furthermore, for $x \in [0, 1]$, $a(x) = 1 \iff x = 1$, implying that, for any $x \in (0, 1)$, $a(x) < 1$, which concludes the proof. \square

proposition B.4 implies that, for any $w \in (0, 1)$, variance reduction can always be achieved provided that $f(\mathbf{X})$ and $g(\mathbf{X})$ are sufficiently correlated.

Proposition B.5. For any $w \in (0, 1)$, $4w/(1 + w)^2 > 2w/(1 + w)$.

Proof. Taking the ratio of the left-hand side to the right-hand side, we have

$$\frac{4w/(1 + w)^2}{2w/(1 + w)} = \frac{2}{1 + w} > 1, \quad (\text{B.6})$$

since $w < 1$. \square

Corollary B.6. If $\beta^{\text{MFMC}} < 1$, then $N' > N$.

Proof. From proposition B.3, if $\beta^{\text{MFMC}} < 1$, then $\rho^2 > 4w/(1 + w)^2$. Then, from proposition B.5, we deduce that $\rho^2 > 2w/(1 + w) > w/(1 + w)$. Finally, we conclude from proposition B.2 that $\eta > 1$, i.e., $N' > N$. \square

B.2 The exact CV estimator

We now turn to the exact CV estimator of the expected value of $f(\mathbf{X})$, i.e., considering that $\tau := \mathbb{E}[g(\mathbf{X})]$ is known. For a given (expected) computational budget \mathcal{C} , the (optimal) CV estimator reads

$$\hat{\mu}_{\mathcal{C}}^{\text{CV}} = \frac{1}{n} \sum_{i=1}^n f(\mathbf{X}^{(i)}) - \alpha \left(\frac{1}{n} \sum_{i=1}^n g(\mathbf{X}^{(i)}) - \tau \right), \quad \text{with } \alpha = \frac{\mathbb{C}[Y, Z]}{\mathbb{V}[Z]}, \quad n = \frac{\mathcal{C}}{c_f(1+w)}, \quad (\text{B.7})$$

and where $\mathcal{X} := \{\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(n)}\}$ is an n -sample of \mathbf{X} . We have the following properties.

Proposition B.7. *The variance of the CV estimator eq. (B.7) is such that $\mathbb{V}[\hat{\mu}_{\mathcal{C}}^{\text{CV}}] = \beta^{\text{CV}} \mathbb{V}[\hat{\mu}_p]$, with $\beta^{\text{CV}} = (1+w)(1-\rho^2)$.*

Proof. From eq. (2.5) and appendix A.1, we know that $\mathbb{V}[\hat{\mu}_{\mathcal{C}}^{\text{CV}}] = (1-\rho^2)\mathbb{V}[\hat{\mu}_n]$. The result follows from the fact that $\mathbb{V}[\hat{\mu}_n] = \mathbb{V}[f(\mathbf{X})]/n$ and $\mathbb{V}[\hat{\mu}_p] = \mathbb{V}[f(\mathbf{X})]/p$, with $n = p/(1+w)$. \square

Proposition B.8. *$\beta^{\text{CV}} < 1$ if and only if $\rho^2 > w/(1+w)$.*

Proof. $\beta^{\text{CV}} < 1 \iff (1+w)(\rho^2 - 1) > -1 \iff \rho^2 > 1 - 1/(1+w) = w/(1+w)$. \square

Combined with propositions B.3 and B.8, proposition B.5 indicates that the minimum value of ρ^2 that guarantees variance reduction is always (i.e., for any $w \in (0, 1)$) at least twice as large for the MFMC estimator than for the CV estimator. For instance, for $w = 1/3$, the MFMC estimator requires $\rho^2 > 3/4$ to have a reduced variance, while the CV estimator only requires $\rho^2 > 1/4$.

Proposition B.9. *For $w, \rho^2 \in (0, 1)$, $\beta^{\text{CV}} < 1$ if and only if $(1+w)\beta^{\text{CV}} < \beta^{\text{MFMC}}$.*

Proof. $\beta^{\text{CV}} < 1 \iff \rho^2 > w/(1+w) \iff 1 - \rho^2 < 1/(1+w) \iff 1/(1-\rho^2) > 1+w$. Because $1/(1-\rho^2) - 1 = \rho^2/(1-\rho^2)$, $\beta^{\text{CV}} < 1 \iff \rho^2/(1-\rho^2) > w \iff w\rho^2/(1-\rho^2) > w^2$. We note that $w\rho^2/(1-\rho^2) = (\eta w)^2$, so that $\beta^{\text{CV}} < 1 \iff \eta w > w$. Therefore, because $\beta^{\text{MFMC}}/\beta^{\text{CV}} = (1+\eta w)^2/(1+w)$, $\beta^{\text{CV}} < 1 \iff \beta^{\text{MFMC}}/\beta^{\text{CV}} > (1+w)^2/(1+w) = 1+w$, which concludes the proof. \square

Corollary B.10. *For $w, \rho^2 \in (0, 1)$, if $\beta^{\text{CV}} < 1$, then $\beta^{\text{CV}} < \beta^{\text{MFMC}}$.*

Proof. Because $w > 0$, $\beta^{\text{MFMC}}/\beta^{\text{CV}} > 1+w \implies \beta^{\text{MFMC}}/\beta^{\text{CV}} > w$. Hence the result follows from proposition B.9. \square

Corollary B.10 implies that, in the (w, ρ^2) variance reduction region of the CV estimator, which is larger than that of the MFMC estimator by propositions B.3, B.5, and B.8, the variance reduction of the CV estimator is always greater than that of the MFMC estimator (i.e., $\beta^{\text{CV}} < \beta^{\text{MFMC}}$).

Figure 5 depicts the (w, ρ^2) variance reduction regions for the CV (fig. 5a) and MFMC (fig. 5b) estimators. We clearly see that for any $w \in (0, 1)$, a higher value of ρ^2 is required for the MFMC estimator to achieve variance reduction than for the CV estimator (red solid line). Figure 5 also illustrates an immediate consequence of propositions B.1 and B.7, namely that $\lim_{w \rightarrow 0} \beta^{\text{CV}} = \lim_{w \rightarrow 0} \beta^{\text{MFMC}} = 1 - \rho^2$. We also observe that β^{CV} converges to $1 - \rho^2$ faster than β^{MFMC} , consistent with corollary B.10.

These properties illustrate the advantages of using surrogate-based control variates. Indeed, when the low-fidelity model g corresponds to a surrogate model, not only $w \ll 1$, but the exact expectation $\mathbb{E}[g(\mathbf{X})]$ may be known, so that the exact CV approach can be used, resulting in a variance reduction factor close to ρ^2 .




Figure 5: Contour lines of β^{CV} (fig. 5a) and β^{MFMC} (fig. 5b) as functions of w and ρ^2 . In each plot, the green region corresponds to pairs (w, ρ^2) for which the variance is reduced, while the red region corresponds to pairs for which the variance is increased. The boundary between these two regions is represented by a red line.

C Optimal MLMC-MLCV variance estimator

The MLMC-MLCV variance estimator is given by eq. (3.20), with

$$\hat{T}_0^{\text{MLCV}}(\alpha_0) = \hat{V}^{(0)}[Y_0] - \alpha_0^{\text{T}}(\hat{V}^{(0)}[\mathbf{Z}] - \sigma_{\mathbf{Z}}^2), \quad (\text{C.1})$$

$$\hat{T}_\ell^{\text{MLCV}}(\alpha_\ell) = \hat{V}^{(\ell)}[Y_\ell] - \hat{V}^{(\ell)}[Y_{\ell-1}] - \alpha_\ell^{\text{T}}[\hat{V}^{(\ell)}[\mathbf{Z}_{1:}] - \hat{V}^{(\ell)}[\tilde{\mathbf{Z}}] - (\sigma_{\mathbf{Z}_{1:}}^2 - \sigma_{\tilde{\mathbf{Z}}}^2)], \text{ for } \ell > 1, \quad (\text{C.2})$$

where

$$\mathbf{Z} = (Z_0, \dots, Z_L) = (g_0(\mathbf{X}), \dots, g_L(\mathbf{X})), \quad \sigma_{\mathbf{Z}}^2 = \mathbb{V}[\mathbf{Z}], \quad (\text{C.3})$$

$$\mathbf{Z}_{1:} = (Z_1, \dots, Z_L) = (g_1(\mathbf{X}), \dots, g_L(\mathbf{X})), \quad \sigma_{\mathbf{Z}_{1:}}^2 = \mathbb{V}[\mathbf{Z}_{1:}], \quad (\text{C.4})$$

$$\tilde{\mathbf{Z}} = (\tilde{Z}_0, \dots, \tilde{Z}_{L-1}) = (\tilde{g}_0(\mathbf{X}), \dots, \tilde{g}_{L-1}(\mathbf{X})), \quad \sigma_{\tilde{\mathbf{Z}}}^2 = \mathbb{V}[\tilde{\mathbf{Z}}], \quad (\text{C.5})$$

$$\tilde{g}_{\ell-1} = g_\ell - h_\ell, \quad \text{for } \ell > 0. \quad (\text{C.6})$$

Furthermore, we assume that the expected values of the control variates are known, so that we use the following variance estimators:

$$\hat{V}^{(0)}[\mathbf{Z}] = \hat{E}^{(0)}[\bar{\mathbf{Z}}^{\odot 2}], \quad \hat{V}^{(\ell)}[\mathbf{Z}_{1:}] - \hat{V}^{(\ell)}[\tilde{\mathbf{Z}}] = \hat{E}^{(\ell)}[\bar{\mathbf{Z}}_{1:}^{\odot 2}] - \hat{E}^{(\ell)}[\tilde{\mathbf{Z}}^{\odot 2}] = \hat{E}^{(\ell)}[\bar{\mathbf{Z}}_{1:}^{\odot 2} - \tilde{\mathbf{Z}}^{\odot 2}]. \quad (\text{C.7})$$

The optimal values α_ℓ^* of the CV parameters are given by

$$\alpha_0^* = \Sigma_0^{-1} \mathbf{c}_0, \quad \Sigma_0 = \mathbb{C}[\bar{\mathbf{Z}}^{\odot 2}], \quad \mathbf{c}_0 = \mathbb{C}[\bar{\mathbf{Z}}^{\odot 2}, \bar{Y}_0^2], \quad (\text{C.8})$$

$$\alpha_\ell^* = \Sigma_\ell^{-1} \mathbf{c}_\ell, \quad \Sigma_\ell = \mathbb{C}[\bar{\mathbf{Z}}_{1:}^{\odot 2} - \tilde{\mathbf{Z}}^{\odot 2}], \quad \mathbf{c}_\ell = \mathbb{C}[\bar{\mathbf{Z}}_{1:}^{\odot 2} - \tilde{\mathbf{Z}}^{\odot 2}, \bar{Y}_\ell^2 - \bar{Y}_{\ell-1}^2], \quad \text{for } \ell > 0. \quad (\text{C.9})$$

Note that Σ_ℓ is the same for all $\ell > 0$. For PC-based control variates, $g_\ell(\mathbf{X}) = \sum_{k=0}^{P_g^\ell} \mathfrak{g}_{\ell,k} \Psi_k(\mathbf{X})$ and $h_\ell(\mathbf{X}) = \sum_{k=0}^{P_h^\ell} \mathfrak{h}_{\ell,k} \Psi_k(\mathbf{X})$, letting $P_g^\ell := \max(P_g^\ell, P_h^\ell)$, we have

$$\tilde{Z}_{\ell-1} = \sum_{k=0}^{P_g^\ell} \tilde{\mathfrak{g}}_{\ell-1,k} \Psi_k(\mathbf{X}), \quad \text{with } \tilde{\mathfrak{g}}_{\ell-1,k} := \mathfrak{g}_{\ell,k} - \mathfrak{h}_{\ell,k}, \quad \text{for } \ell = 1, \dots, L, \quad (\text{C.10})$$

where $\mathfrak{g}_{\ell,k} := 0$ for $k > P_g^\ell$ and $\mathfrak{h}_{\ell,k} := 0$ for $k > P_h^\ell$. Consequently,

$$\mathbb{E}[Z_\ell] = \mathfrak{g}_{\ell,0}, \quad \mathbb{E}[\tilde{Z}_{\ell-1}] = \tilde{\mathfrak{g}}_{\ell-1,0}, \quad (\text{C.11})$$

$$\mathbb{V}[Z_\ell] = \sum_{k=1}^{P_g^\ell} \mathfrak{g}_{\ell,k}^2, \quad \mathbb{V}[\tilde{Z}_{\ell-1}] = \sum_{k=1}^{P_g^\ell} \tilde{\mathfrak{g}}_{\ell-1,k}^2. \quad (\text{C.12})$$

Furthermore,

$$[\Sigma_0]_{m,m'} = \sum_{i,j=1}^{P_g^m} \sum_{q,r=1}^{P_g^{m'}} \mathfrak{g}_{m,i} \mathfrak{g}_{m,j} \mathfrak{g}_{m',q} \mathfrak{g}_{m',r} (\Phi_{ijqr} - \delta_{ij} \delta_{qr}), \quad \text{for } m, m' = 0, \dots, L, \quad (\text{C.13})$$

where $\Phi_{ijqr} := \mathbb{E}[\Psi_i(\mathbf{X}) \Psi_j(\mathbf{X}) \Psi_q(\mathbf{X}) \Psi_r(\mathbf{X})]$ are the entries of the fourth-order Galerkin product tensor, which is a well-known object in stochastic Galerkin methods [20, 32], and δ_{ij} denotes the Kronecker delta. Besides, noticing that $\tilde{\mathbf{Z}}_1^{\odot 2} - \tilde{\mathbf{Z}}^{\odot 2} = (\tilde{\mathbf{Z}}_1 - \tilde{\mathbf{Z}}) \odot (\tilde{\mathbf{Z}}_1 + \tilde{\mathbf{Z}}) = \tilde{\mathbf{W}} \odot (\tilde{\mathbf{Z}}_1 + \tilde{\mathbf{Z}})$, with $\tilde{\mathbf{W}} = \mathbf{W} - \boldsymbol{\mu}_{\mathbf{W}}$ and the definitions in eq. (3.30), we have

$$[\Sigma_\ell]_{m,m'} = A_{m,m'} + B_{m,m'} + C_{m,m'} + C_{m',m}, \quad \text{for } \ell, m, m' = 1, \dots, L, \quad (\text{C.14})$$

where

$$A_{m,m'} = \mathbb{C}[\bar{W}_m \bar{Z}_m, \bar{W}_{m'} \bar{Z}_{m'}] \quad (\text{C.15})$$

$$= \sum_{i=1}^{P_h^m} \sum_{j=1}^{P_g^m} \sum_{q=1}^{P_h^{m'}} \sum_{r=1}^{P_g^{m'}} \mathfrak{h}_{m,i} \mathfrak{g}_{m,j} \mathfrak{h}_{m',q} \mathfrak{g}_{m',r} (\Phi_{ijqr} - \delta_{ij} \delta_{qr}), \quad (\text{C.16})$$

$$B_{m,m'} = \mathbb{C}[\bar{W}_m \tilde{\tilde{Z}}_{m-1}, \bar{W}_{m'} \tilde{\tilde{Z}}_{m'-1}] \quad (\text{C.17})$$

$$= \sum_{i=1}^{P_h^m} \sum_{j=1}^{P_g^m} \sum_{q=1}^{P_h^{m'}} \sum_{r=1}^{P_g^{m'}} \mathfrak{h}_{m,i} \tilde{\mathfrak{g}}_{m-1,j} \mathfrak{h}_{m',q} \tilde{\mathfrak{g}}_{m'-1,r} (\Phi_{ijqr} - \delta_{ij} \delta_{qr}), \quad (\text{C.18})$$

$$C_{m,m'} = \mathbb{C}[\bar{W}_m \tilde{\tilde{Z}}_{m-1}, \bar{W}_{m'} \bar{Z}_{m'}] \quad (\text{C.19})$$

$$= \sum_{i=1}^{P_h^m} \sum_{j=1}^{P_g^m} \sum_{q=1}^{P_h^{m'}} \sum_{r=1}^{P_g^{m'}} \mathfrak{h}_{m,i} \tilde{\mathfrak{g}}_{m-1,j} \mathfrak{h}_{m',q} \mathfrak{g}_{m',r} (\Phi_{ijqr} - \delta_{ij} \delta_{qr}). \quad (\text{C.20})$$

D Taylor surrogate for the numerical test case

In our example, f_ℓ is only differentiable in $[-\pi, \pi]^3 \times [\nu_{\min}, \nu_{\max}] \times ([-1, 0) \cup (0, 1])^3$. Consequently, the Taylor polynomial surrogate cannot be used directly as defined in (2.17) and (2.18), because

the Jacobian and Hessian matrices are not defined at $\boldsymbol{\mu}_{\mathbf{X}}$. Instead, we define the first-order Taylor surrogate as

$$f_{\ell}(\mathbf{X}) \simeq g_{\ell}^{\text{T}_1}(\mathbf{X}) = f_{\ell}(\boldsymbol{\mu}_{\mathbf{X}}) + \sum_{i=1}^7 g_{\ell,i}^{\text{T}_1}(\boldsymbol{\mu}_{\mathbf{X}}; \mathbf{X}), \quad (\text{D.1})$$

where, for $i = 1, \dots, 4$, $g_{\ell,i}^{\text{T}_1}(\boldsymbol{\mu}_{\mathbf{X}}; \mathbf{X}) = (X_i - \mu_{X_i}) \frac{\partial f_{\ell}}{\partial X_i}(\boldsymbol{\mu}_{\mathbf{X}})$, and, for $i = 5, \dots, 7$,

$$g_{\ell,i}^{\text{T}_1}(\boldsymbol{\mu}_{\mathbf{X}}; \mathbf{X}) = (X_i - \mu_{X_i}) \times \begin{cases} \frac{\partial f_{\ell}}{\partial X_i}(\boldsymbol{\mu}_{\mathbf{X}}), & \mu_{X_i} \neq 0, \\ \lim_{\boldsymbol{\mu}'_{\mathbf{X}} \rightarrow \boldsymbol{\mu}_{\mathbf{X}}^{i,0^-}} \frac{\partial f_{\ell}}{\partial X_i}(\boldsymbol{\mu}'_{\mathbf{X}}), & \mu_{X_i} = 0, X_i < 0, \\ \lim_{\boldsymbol{\mu}'_{\mathbf{X}} \rightarrow \boldsymbol{\mu}_{\mathbf{X}}^{i,0^+}} \frac{\partial f_{\ell}}{\partial X_i}(\boldsymbol{\mu}'_{\mathbf{X}}), & \mu_{X_i} = 0, X_i > 0, \\ 0 & \mu_{X_i} = X_i = 0, \end{cases} \quad (\text{D.2})$$

where $\boldsymbol{\mu}_{\mathbf{X}}^{i,0^{\pm}} = (\mu_{X_1}, \dots, \mu_{X_{i-1}}, 0^{\pm}, \mu_{X_{i+1}}, \dots, \mu_{X_7})$, which is now well-defined.

With our choice of distributions for \mathbf{X} given in (4.11), we have $\boldsymbol{\mu}_{\mathbf{X}} = (0, 0, 0, 0.005, 0, 0, 0)$, and the first-order Taylor surrogate is defined by (D.1), with

$$g_{\ell,1}^{\text{T}_1}(\boldsymbol{\mu}_{\mathbf{X}}; \mathbf{X}) = 7X_1 \sum_{k=1}^K B_k^{\ell}(T; \boldsymbol{\mu}_{\mathbf{X}}) \sum_{j=1}^{N_{\ell}} w_j \sin(k\pi x_j) \mathcal{F}_2(x_j), \quad (\text{D.3})$$

$$g_{\ell,2}^{\text{T}_1}(\boldsymbol{\mu}_{\mathbf{X}}; \mathbf{X}) = 0, \quad (\text{D.4})$$

$$g_{\ell,3}^{\text{T}_1}(\boldsymbol{\mu}_{\mathbf{X}}; \mathbf{X}) = 0, \quad (\text{D.5})$$

$$g_{\ell,4}^{\text{T}_1}(\boldsymbol{\mu}_{\mathbf{X}}; \mathbf{X}) = -(X_4 - 0.005)\pi^2 T \sum_{k=1}^K k^2 \exp(-0.005k^2\pi^2 T) A_k^{\ell}(\boldsymbol{\mu}_{\mathbf{X}}) \sum_{j=1}^{N_{\ell}} w_j \sin(k\pi x_j), \quad (\text{D.6})$$

$$g_{\ell,5}^{\text{T}_1}(\boldsymbol{\mu}_{\mathbf{X}}; \mathbf{X}) = 400|X_5| \sum_{k=1}^K B_k^{\ell}(T; \boldsymbol{\mu}_{\mathbf{X}}) \sum_{j=1}^{N_{\ell}} w_j \sin(k\pi x_j) \mathcal{F}_1(x_j), \quad (\text{D.7})$$

$$g_{\ell,6}^{\text{T}_1}(\boldsymbol{\mu}_{\mathbf{X}}; \mathbf{X}) = 400|X_6| \sum_{k=1}^K B_k^{\ell}(T; \boldsymbol{\mu}_{\mathbf{X}}) \sum_{j=1}^{N_{\ell}} w_j \sin(k\pi x_j) \mathcal{F}_1(x_j), \quad (\text{D.8})$$

$$g_{\ell,7}^{\text{T}_1}(\boldsymbol{\mu}_{\mathbf{X}}; \mathbf{X}) = 400|X_7| \sum_{k=1}^K B_k^{\ell}(T; \boldsymbol{\mu}_{\mathbf{X}}) \sum_{j=1}^{N_{\ell}} w_j \sin(k\pi x_j) \mathcal{F}_1(x_j). \quad (\text{D.9})$$

Because of the piecewise definition of $g_{\ell}^{\text{T}_1}$, the identity $\mathbb{E}[g_{\ell}^{\text{T}_1}(\mathbf{X})] = f_{\ell}(\boldsymbol{\mu}_{\mathbf{X}})$ for a regular first-order Taylor surrogate no longer holds. Instead, we have

$$\mathbb{E}[g_{\ell}^{\text{T}_1}(\mathbf{X})] = f_{\ell}(\boldsymbol{\mu}_{\mathbf{X}}) + 600 \sum_{k=1}^K B_k^{\ell}(T; \boldsymbol{\mu}_{\mathbf{X}}) \sum_{j=1}^{N_{\ell}} w_j \sin(k\pi x_j) \mathcal{F}_1(x_j). \quad (\text{D.10})$$

References

- [1] T. W. Anderson. *An Introduction to Multivariate Statistical Analysis*. Wiley Series in Probability and Statistics. Wiley-Interscience, Hoboken, N.J, 3rd ed edition, 2003.

- [2] Ivo Babuška, Fabio Nobile, and Raúl Tempone. A Stochastic Collocation Method for Elliptic Partial Differential Equations with Random Input Data. *SIAM Journal on Numerical Analysis*, 45(3):1005–1034, 2007.
- [3] Marc Berveiller, Bruno Sudret, and Maurice Lemaire. Stochastic finite element: A non intrusive approach by regression. *European Journal of Computational Mechanics*, 15(1-3):81–92, 2006.
- [4] Claudio Bierig and Alexey Chernov. Convergence analysis of multilevel Monte Carlo variance estimators and application for random obstacle problems. *Numerische Mathematik*, 130(4):579–613, 2015.
- [5] Claudio Bierig and Alexey Chernov. Estimation of arbitrary order central statistical moments by the multilevel Monte Carlo method. *Stochastics and Partial Differential Equations Analysis and Computations*, 4(1):3–40, 2016.
- [6] Géraud Blatman and Bruno Sudret. Efficient computation of global sensitivity indices using sparse polynomial chaos expansions. *Reliability Engineering & System Safety*, 95(11):1216–1229, 2010.
- [7] Géraud Blatman and Bruno Sudret. Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, 230(6):2345–2367, 2011.
- [8] François-Xavier Briol, Chris J. Oates, Mark Girolami, Michael A. Osborne, and Dino Sejdinovic. Probabilistic Integration: A Role in Statistical Computation? *Statistical Science*, 34(1):1–22, 2019.
- [9] Robert H. Cameron and William T. Martin. The orthogonal development of non-linear functionals in series of Fourier-Hermite functionals. *Annals of Mathematics*, 48(2):385–392, 1947.
- [10] K. A. Cliffe, M. B. Giles, R. Scheichl, and A. L. Teckentrup. Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Computing and Visualization in Science*, 14(1):3–15, 2011.
- [11] Patrick R. Conrad and Youssef M. Marzouk. Adaptive Smolyak Pseudospectral Approximations. *SIAM Journal on Scientific Computing*, 35(6):A2643–A2670, 2013.
- [12] Paul G. Constantine, Michael S. Eldred, and Eric T. Phipps. Sparse pseudospectral approximation method. *Computer Methods in Applied Mechanics and Engineering*, 229–232:1–12, 2012.
- [13] Josef Dick, Frances Y. Kuo, and Ian H. Sloan. High-dimensional integration: The quasi-monte carlo way. *Acta Numerica*, 22:133–288, 2013.
- [14] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least Angle Regression. *The Annals of Statistics*, 32(2):407–451, 2004.
- [15] Reda EL Amri, Rodolphe Le Riche, Céline Helbert, Christophette Blanchet-Scalliet, and Sébastien Da Veiga. A sampling criterion for constrained bayesian optimization with uncertainties. *arXiv preprint arXiv:2103.05706*, 2021.
- [16] Jamie Fox and Giray Ökten. Polynomial chaos as a control variate method. *SIAM Journal on Scientific Computing*, 43(3):A2268–A2294, 2021.

- [17] Sarthak Garg, Neeraj Sood, and Costas D. Sarris. Uncertainty quantification of ray-tracing based wireless propagation models with a control variate-polynomial chaos expansion method. In *2015 IEEE International Symposium on Antennas and Propagation & USNC/URSI National Radio Science Meeting*, pages 1776–1777, 2015.
- [18] G Geraci, M Eldred, and G Iaccarino. A multifidelity control variate approach for the multilevel Monte Carlo technique. Technical report, Stanford University, 2015.
- [19] Gianluca Geraci, Michael S. Eldred, and Gianluca Iaccarino. A multifidelity multilevel Monte Carlo method for uncertainty propagation in aerospace applications. In *19th AIAA Non-Deterministic Approaches Conference*, Grapevine, Texas, 2017. American Institute of Aeronautics and Astronautics.
- [20] Roger G. Ghanem and Pol D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Courier Corporation, 2003.
- [21] Michael B. Giles. Multilevel Monte Carlo Path Simulation. *Operations Research*, 56(3):607–617, 2008.
- [22] Michael B. Giles. Multilevel Monte Carlo methods. *Acta Numerica*, 24:259–328, 2015.
- [23] Michael B. Giles and Benjamin J. Waterhouse. Multilevel quasi-Monte Carlo path simulation. *Advanced Financial Modelling, Radon Series on Computational and Applied Mathematics*, 8:165–181, 2009.
- [24] Alex A. Gorodetsky, Gianluca Geraci, Michael S. Eldred, and John D. Jakeman. A generalized approximate control variate framework for multifidelity uncertainty quantification. *Journal of Computational Physics*, 408:109257, 2020.
- [25] Zixi Gu and Costas D. Sarris. Multi-parametric uncertainty quantification with a hybrid monte-carlo / polynomial chaos expansion ftd method. In *2015 IEEE MTT-S International Microwave Symposium*, pages 1–3, 2015.
- [26] Abdul-Lateef Haji-Ali, Fabio Nobile, and Raúl Tempone. Multi-index Monte Carlo: When sparsity meets sampling. *Numerische Mathematik*, 132(4):767–806, 2016.
- [27] Tim Hesterberg. Weighted average importance sampling and defensive mixture distributions. *Technometrics*, 37(2):185–194, 1995.
- [28] Janis Janusevskis and Rodolphe Le Riche. Simultaneous kriging-based estimation and optimization of mean response. *Journal of Global Optimization*, 55(2):313–336, 2013.
- [29] S. S Lavenberg, T. L. Moeller, and P. D. Welch. The application of control variables to the simulation of closed queueing networks. In *Proceedings of the 9th Conference on Winter Simulation - Volume 1, WSC '77*, pages 152–154, Gaithersburg, Maryland, USA, 1977. Winter Simulation Conference.
- [30] S. S. Lavenberg and P. D. Welch. A Perspective on the Use of Control Variables to Increase the Efficiency of Monte Carlo Simulations. *Management Science*, 27(3):322–335, 1981.
- [31] Stephen S. Lavenberg, Thomas L. Moeller, and Peter D. Welch. Statistical Results on Control Variables with Application to Queueing Network Simulation. *Operations Research*, 30(1):182–202, 1982.

- [32] O. P. Le Maître and Omar M. Knio. *Spectral Methods for Uncertainty Quantification*. Scientific Computation. Springer Netherlands, Dordrecht, 2010.
- [33] Christiane Lemieux. Control Variates. In *Wiley StatsRef: Statistics Reference Online*, pages 1–8. American Cancer Society, 2017.
- [34] Nora Lüthen, Stefano Marelli, and Bruno Sudret. Sparse polynomial chaos expansions: Literature survey and benchmark. *SIAM/ASA Journal on Uncertainty Quantification*, 9(2):593–649, 2021.
- [35] MD McKay, RJ Beckman, and WJ Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245, 1979.
- [36] Paul Mycek and Matthias De Lozzo. Multilevel Monte Carlo Covariance Estimation for the Computation of Sobol’ Indices. *SIAM/ASA Journal on Uncertainty Quantification*, 7(4):1323–1348, 2019.
- [37] Habib N. Najm. Uncertainty quantification and polynomial chaos techniques in computational fluid dynamics. *Annual Review of Fluid Mechanics*, 41(1):35–52, 2009.
- [38] Barry L. Nelson. Control Variate Remedies. *Operations Research*, 38(6):974–992, 1990.
- [39] Leo W. T. Ng and Karen E. Willcox. Multifidelity approaches for optimization under uncertainty. *International Journal for Numerical Methods in Engineering*, 100(10):746–772, 2014.
- [40] F. Nobile, R. Tempone, and C. G. Webster. An Anisotropic Sparse Grid Stochastic Collocation Method for Partial Differential Equations with Random Input Data. *SIAM Journal on Numerical Analysis*, 46(5):2411–2442, 2008.
- [41] J. Park and I. W. Sandberg. Universal Approximation Using Radial-Basis-Function Networks. *Neural Computation*, 3(2):246–257, 06 1991.
- [42] Raghu Pasupathy, Bruce W. Schmeiser, Michael R. Taaffe, and Jin Wang. Control-variate estimation using estimated control means. *IIE Transactions*, 44(5):381–385, 2012.
- [43] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Optimal Model Management for Multifidelity Monte Carlo Estimation. *SIAM Journal on Scientific Computing*, 38(5):A3163–A3194, 2016.
- [44] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization. *SIAM Review*, 60(3):550–591, 2018.
- [45] Carl Edward Rasmussen. *Gaussian Processes in Machine Learning*, pages 63–71. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
- [46] Matthew T. Reagan, Habib N. Najm, Roger G. Ghanem, and Omar M. Knio. Uncertainty quantification in reacting-flow simulations through non-intrusive spectral projection. *Combustion and Flame*, 132(3):545–555, 2003.
- [47] Daniel Schaden. *Variance Reduction with Multilevel Estimators*. PhD thesis, Technische Universität München, Munich, 2021.

- [48] Daniel Schaden and Elisabeth Ullmann. On Multilevel Best Linear Unbiased Estimators. *SIAM/ASA Journal on Uncertainty Quantification*, 8(2):601–635, 2020.
- [49] Daniel Schaden and Elisabeth Ullmann. Asymptotic Analysis of Multilevel Best Linear Unbiased Estimators. *SIAM/ASA Journal on Uncertainty Quantification*, 9(3):953–978, 2021.
- [50] Ralph C. Smith. *Uncertainty Quantification: Theory, Implementation, and Applications*. Society for Industrial and Applied Mathematics, USA, 2013.
- [51] Bruno Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, 93(7):964–979, 2008.
- [52] T. J. Sullivan. *Introduction to Uncertainty Quantification*, volume 63. Springer Cham, 2015.
- [53] Aretha Leonore Teckentrup. *Multilevel Monte Carlo Methods and Uncertainty Quantification*. PhD thesis, University of Bath, 2013.
- [54] Rohit K. Tripathy and Ilias Bilionis. Deep uq: Learning deep neural network surrogate models for high dimensional uncertainty quantification. *Journal of Computational Physics*, 375:565–588, 2018.
- [55] Norbert Wiener. The Homogeneous Chaos. *American Journal of Mathematics*, 60(4):897–936, 1938.
- [56] Hang Yang, Yuji Fujii, K. W. Wang, and Alex A. Gorodetsky. Control Variate Polynomial Chaos: Optimal Fusion of Sampling and Surrogates for Multifidelity Uncertainty Quantification. *arXiv:2201.10745 [stat]*, 2022.