

Multi-output multilevel best linear unbiased estimators via semidefinite programming

M. Croci^{a,*}, K. E. Willcox^a, S. J. Wright^b

^a*Oden Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, TX, USA.*

^b*Computer Science Department, University of Wisconsin-Madison, Madison, WI, USA.*

Abstract

Multifidelity forward uncertainty quantification (UQ) problems often involve multiple quantities of interest and heterogeneous models (e.g., different grids, equations, dimensions, physics, surrogate and reduced-order models). While computational efficiency is key in this context, multi-output strategies in multilevel/multifidelity methods are either sub-optimal or non-existent. In this paper we extend multilevel best linear unbiased estimators (MLBLUE) to multi-output forward UQ problems and we present new semidefinite programming formulations for their optimal setup. Not only do these formulations yield the optimal number of samples required, but also the optimal selection of low-fidelity models to use. While existing MLBLUE approaches are single-output only and require a non-trivial nonlinear optimization procedure, the new multi-output formulations can be solved reliably and efficiently. We demonstrate the efficacy of the new methods and formulations in practical UQ problems with model heterogeneity.

Key words: Multilevel Monte Carlo, multifidelity Monte Carlo, semidefinite programming, uncertainty quantification, sample allocation, model selection.

1. Introduction

A central task in forward uncertainty quantification (UQ) is estimating the expectation of output quantities of interest (QoI). When the same QoI is the output of different models of varying fidelity, multilevel and multifidelity Monte Carlo algorithms [1, 2, 3, 4, 5, 6, 7, 8, 9] efficiently estimate its expectation via sampling: by exploiting the correlations between models, these algorithms reduce the estimator variance for a fixed cost (or, equivalently, reduce the total cost to obtain an estimator with a given level of variance). Multilevel and multifidelity methods require selecting which models to sample and the number of samples to be drawn, a choice that can impact efficiency if sub-optimal. But setting up these methods optimally is a non-trivial operation, especially when UQ problems involve multiple QoIs and heterogeneous models, as is often the case in many engineering and scientific applications. We say models are “heterogeneous” if they are constructed through different approaches (e.g., varying grids, equations, dimensions, or physics, or using surrogate and reduced-order models) or if their output structure varies (e.g., if models have different QoIs). Since it is crucial to design estimators that are as efficient as possible, in this paper we extend multilevel best linear unbiased estimators (MLBLUEs) [8, 9] to the multi-output case and propose a fast, reliable strategy for their optimal model selection and sample allocation.

There are many multilevel and multifidelity estimators available in the literature [1, 2, 3, 4, 5, 6, 7, 8, 9], each differing in the way they exploit the statistical relations between models. A key ingredient of any multilevel/multifidelity algorithm is a strategy for optimal model selection and sample allocation: given a set of available models with different characteristics, it is important to determine which models actually lead to a performance benefit, and how much these models should be sampled. The optimal *model selection and sample allocation problem* (MOSAP) depends on both the method and on the set of available models.

*Corresponding author

In some specific cases the MOSAP is trivial to solve; the model selection largely simplifies when all models belong to the same discretization hierarchy [2, 4]. However, in the general case in which models are highly heterogeneous, the MOSAP is an NP-hard nonlinear mixed-integer programming problem, and its complexity increases exponentially with the number of available models.

In this paper we focus on MBLUES, and specifically on techniques for solving their MOSAP. MBLUES were presented by Schaden and Ullmann in [8, 9], and have the appealing property of being provably optimal among all multilevel linear unbiased estimators¹ [8, 9]. Furthermore, while multilevel Monte Carlo (MLMC) [1, 2, 4] can only couple² models in pairs, and multifidelity Monte Carlo (MFMC) [3, 6] requires all models to be coupled together, MBLUES have no coupling restrictions and can work with any combination of models [8].

Nevertheless, setting up a MBLUE presents challenges. The MOSAP formulation in [8] is a (weakly) convex nonlinear problem with linear constraints, but its objective function is defined via the inverse of a matrix that may be singular. This is handled by adding a small positive term to the diagonal of the matrix, but the resulting problem is still not strongly convex in general and may be ill-conditioned, making the MBLUE MOSAP a hard problem to solve. For instance, the generic optimization solver `fmincon` is used in the numerical tests of [8, Section 6], but it fails in several cases. Another limitation of the existing MBLUES is that they have been designed only for a single output. Extending these estimators to multiple QoIs, while trivial in theory, can be very challenging in practice since any difficulties in the MBLUE MOSAP optimization are likely to be exacerbated in the multi-output case.

In this paper we make the following contributions.

- We reformulate the MBLUE MOSAP as a semidefinite programming (SDP) problem. SDPs and linear programming are both conic optimization problems, a special class of convex optimization [10]. SDPs have been an area of intense study in the optimization community since the mid-1990s, and powerful algorithms and software have been developed [10, 11, 12, 13, 14]. Our SDP reformulation does not require taking the inverse of a possibly ill-conditioned matrix, nor does it require a positive shift to ensure positive definiteness of this matrix. We prove the equivalence between the SDP reformulation and the original MOSAP problem with zero shift from [8].
- The original MBLUE MOSAP minimizes the estimator variance for fixed cost [8]. We present a SDP reformulation minimizing the MBLUE cost for a given statistical error tolerance. Depending on the UQ problem, either formulation may be useful, and they can both be solved efficiently.
- We introduce a multi-objective MOSAP SDP formulation that permits computation of Pareto points defined by the two objectives of statistical error and cost. This is the first such formulation in the literature and is of practical relevance in cases where an appropriate limit on the computational budget is not known *a priori*.
- We extend MBLUES to the multi-output case and present an SDP reformulation for the multi-output MOSAP that can be solved efficiently. We show that this multi-output estimator outperforms MLMC and MFMC in numerical experiments.
- We demonstrate numerically that MBLUES maintain their efficiency when the prohibitive cost of high-fidelity models imposes sample restrictions and/or prevents the estimation of all model correlations.

This paper is structured as follows. In Section 2 we introduce MBLUES and their MOSAP and we discuss their properties. We reformulate the single-output MBLUE MOSAP as an SDP in Section 3, where we also present a new formulation in which the total MBLUE cost is minimized subject to a statistical error constraint and a new multi-objective MOSAP formulation. In Section 4 we extend MBLUES and their MOSAPs to the multi-output case. Numerical experiments are presented in Section 5, and Section 6 contains some concluding remarks.

¹This includes most multilevel and multifidelity methods and all previously mentioned methods.

²We denote a group of models to be “coupled” if the algorithm samples all these models with the same input.

Notation. We adopt the following notation:

Scalars, vectors, matrices, sets. We indicate scalars with lower-case letters (e.g., a, b, c), vectors with lower-case bold letters (e.g., $\mathbf{a}, \mathbf{b}, \mathbf{c}$), matrices with upper-case letters (e.g., A, B, C), and sets with calligraphic upper-case letters (e.g., $\mathcal{A}, \mathcal{B}, \mathcal{C}$). Random variables and vectors are indicated in the same way. We indicate with $2^{\mathcal{A}}$ the power set of \mathcal{A} , i.e., the set of all possible subsets of \mathcal{A} .

Linear algebra. Given a matrix A , we indicate with $\text{null}(A)$, $\text{col}(A)$, and $\text{row}(A)$ its nullspace, column and row space respectively. Furthermore, we denote with A^T , A^{-1} and A^\dagger its transpose, inverse, and Moore-Penrose pseudo-inverse, and we write $A \succeq 0$ to indicate that A is positive semi-definite. Let $\{a_i\}_{i=1}^\ell$, $\{\mathbf{a}_i\}_{i=1}^\ell$, and $\{A_i\}_{i=1}^\ell$ be ordered sets of scalars, vectors and matrices respectively. We indicate with $[a_i]_{i=1}^\ell$, $[\mathbf{a}_i]_{i=1}^\ell$, and $[A_i]_{i=1}^\ell$ the vectors and the matrix obtained by vertically stacking the elements (assuming that matrix dimensions are compatible).

Random input parameters. We indicate with ω a generic input parameter which is modeled as random from a known distribution. We indicate independent identically distributed input parameter samples with different subscripts and superscripts, e.g., ω_i, ω_j or ω_i^k, ω_j^s , where the samples are independent if $i \neq j$ or if $k \neq s$.

Random variables, expectation, variance, covariance. We indicate random variables and vectors as functions of the random input ω , e.g. $p(\omega)$ for a scalar random variable or $\mathbf{p}(\omega)$ for a random vector. We indicate with $\mathbb{E}[\cdot]$, $\mathbb{V}[\cdot]$ and $C(\cdot, \cdot)$ the expectation, variance, and covariance operators. We always assume that the random quantities involved have finite variance.

2. Background on multilevel best linear unbiased estimators

To set the scene for the later discussion, we introduce the single-output MLBLUEs by Schaden and Ullmann [8, 9], and discuss their formulation of the MLBLUE MOSAP. This section summarizes the technical approach presented in [8].

In forward UQ the objective is to approximate the expectation $\mathbb{E}[p(\omega)]$ of a scalar QoI $p(\omega)$ which is modeled as a random variable of finite variance. For this purpose, we consider a set of ℓ different models $\mathcal{M} = \{\mathcal{M}_i\}_{i=1}^\ell$ of different accuracy/fidelities, so that for all $i = 1, \dots, \ell$, model \mathcal{M}_i admits the random variable $p_i(\omega)$ as output, where $p_i \approx p$. Here we indicate with \mathcal{M}_1 the highest-fidelity model whose output p_1 best approximates p . Model \mathcal{M}_1 typically has the highest cost across the model set.

Remark 2.1. We present MLBLUEs in a *multifidelity setting* [15] in which p_1 is taken as the best model we have access to. In this setting it is not possible to refine p_1 further. Since the bias error $|\mathbb{E}[p - p_1]|$ cannot be reduced, it is only taken into account when determining a suitable statistical error tolerance and otherwise ignored. This is in contrast to what we call a *multilevel setting* [2] in which it is always possible to refine models further to reduce the bias. We remark that MLBLUEs – as well as MLMC and MFMC – can be used in both settings [8, 9].

Given a set $\mathcal{G} \subseteq 2^{\mathcal{M}} \setminus \emptyset$ of allowed model groupings, MLBLUEs approximate $\mathbb{E}[p_1]$ by combining independent samples of the model groups in \mathcal{G} , where by “sample of a model group” we mean that the models in the group are sampled with the same input (i.e., same ω). We order model groups by size and then lexicographically by model, and we indicate the k -th group in \mathcal{G} with \mathcal{G}_k for $k = 1, \dots, |\mathcal{G}|$. For instance, for $\ell = 3$ and all groupings allowed, we would have $\mathcal{G} = \{\mathcal{G}_k\}_{k=1}^7$, where

$$\begin{aligned} \mathcal{G}_1 &= \{\mathcal{M}_1\}, & \mathcal{G}_2 &= \{\mathcal{M}_2\}, & \mathcal{G}_3 &= \{\mathcal{M}_3\}, \\ \mathcal{G}_4 &= \{\mathcal{M}_1, \mathcal{M}_2\}, & \mathcal{G}_5 &= \{\mathcal{M}_1, \mathcal{M}_3\}, & \mathcal{G}_6 &= \{\mathcal{M}_2, \mathcal{M}_3\}, \\ \mathcal{G}_7 &= \{\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3\}. \end{aligned}$$

Now, let $\boldsymbol{\mu} = [\mathbb{E}[p_i]]_{i=1}^\ell \in \mathbb{R}^\ell$ be the vector of the expectations of all model outputs. Then a sample $\mathbf{p}^k(\omega) = [p_i(\omega)]_{M_i \in \mathcal{G}_k}$ from the k -th group $\mathcal{G}_k \in \mathcal{G}$ satisfies the linear model

$$\mathbf{p}^k(\omega) = \mathbb{E}[\mathbf{p}^k] + (\mathbf{p}^k(\omega) - \mathbb{E}[\mathbf{p}^k]) = R_k \boldsymbol{\mu} + \boldsymbol{\varepsilon}^k(\omega), \quad (1)$$

where $R_k \in \mathbb{R}^{|\mathcal{G}_k| \times \ell}$ is a boolean restriction matrix (made up of rows of the $\ell \times \ell$ identity matrix corresponding to the elements of \mathcal{G}_k) such that $R_k \boldsymbol{\mu} = \mathbb{E}[\mathbf{p}^k]$, and $\boldsymbol{\varepsilon}^k = \mathbf{p}^k(\omega) - \mathbb{E}[\mathbf{p}^k] \in \mathbb{R}^{|\mathcal{G}_k|}$ is a zero-mean random vector with covariance matrix $C_k = C(\mathbf{p}^k, \mathbf{p}^k)$.

The starting point for constructing a MLBLUE is to build on (1) and account for multiple independent samples from all feasible groups. We define the set of group indices $\mathcal{K} := \{1, \dots, |\mathcal{G}|\}$, the number of samples from each group n_k , $k \in \mathcal{K}$, and the sample set for each group to be $\mathcal{J}^k = \{1, \dots, n_k\}$. The vector of these sample sizes — the key variable to be determined in our MOSAP formulation — is $\mathbf{n} := [n_k]_{k \in \mathcal{K}} \in \mathbb{N}^{|\mathcal{G}|}$. Using this notation, we can extend (1) by vertically stacking all samples from all groups into a vector \mathbf{p} such that $\mathbf{p} = [\mathbf{p}^k(\omega_j^k)]_{j \in \mathcal{J}^k, k \in \mathcal{K}}$, where $\mathbf{p}^k(\omega_j^k)$ denotes the j -th sample of model group k and $\mathbf{p}^k(\omega_j^k)$ and $\mathbf{p}^s(\omega_i^s)$ are independent if $s \neq k$ or if $i \neq j$. The vector \mathbf{p} satisfies the linear model

$$\mathbf{p}(\omega) = R\boldsymbol{\mu} + \boldsymbol{\varepsilon}(\omega), \quad (2a)$$

$$\text{where } R = [R_k]_{j \in \mathcal{J}^k, k \in \mathcal{K}}, \quad \boldsymbol{\varepsilon} = [\boldsymbol{\varepsilon}^k(\omega_j^k)]_{j \in \mathcal{J}^k, k \in \mathcal{K}}. \quad (2b)$$

By construction, $\boldsymbol{\varepsilon}$ is a zero-mean vector with block-diagonal covariance since it contains n_k i.i.d. samples of $\boldsymbol{\varepsilon}^k$ for all $k \in \mathcal{K}$ (\mathcal{J}^k is the set of sample indices and $|\mathcal{J}^k| = n_k$). Accordingly, the matrix R contains n_k copies of R_k for each $k \in \mathcal{K}$. Similarly, we define

$$C_\boldsymbol{\varepsilon} := C(\boldsymbol{\varepsilon}, \boldsymbol{\varepsilon}) = \text{diag}(\{C_k\}_{j \in \mathcal{J}^k, k \in \mathcal{K}}), \quad (3)$$

where this block-diagonal matrix contains n_k copies of C_k , for each $k \in \mathcal{K}$. So far \mathbf{n} has been considered fixed. A MLBLUE is found by computing the best linear unbiased estimator (BLUE) of the linear model (2), and then optimizing it with respect to \mathbf{n} .

Remark 2.2. Schaden & Ullmann [8] call the BLUE resulting for any linear model such as (2) that includes multilevel/multifidelity samples a MLBLUE. With the purpose of simplifying the exposition by reducing the number of acronyms, we here call MLBLUE only a *model selection and sample allocation optimal* BLUE for $\boldsymbol{\mu}$ (termed sample allocation optimal MLBLUE or SAOB in [8]).

Note that there is a closed-form expression for the BLUE of (2) independently of whether R and $C_\boldsymbol{\varepsilon}$ are full-rank or not, see e.g., [16]. To proceed, we however make the following simplifying assumption that is also adopted in [8].

Assumption 2.1. The covariance matrices C_k are non-singular for all $k \in \mathcal{K}$.

Remark 2.3. Assumption 2.1 can be removed at the cost of a more complex technical and algorithmic treatment, but is of little importance in practice. Indeed, as shown in [16], a singular C_k implies that there exists a linear combination of outputs from the different models in the set \mathcal{G}_k that is a.s. constant. If this happens, then at least one model is redundant. Once the redundant models are removed from \mathcal{G}_k , the resulting ensemble of models would have a covariance matrix that is full-rank. Thus, Assumption 2.1 is not restrictive.

Remark 2.4. In practice, the matrices C_k often have to be estimated via pilot samples. Since a sample covariance computed with n samples cannot have more than rank n , one must therefore take enough samples to ensure that the sample covariance is not artificially singular. If a singular matrix is nevertheless obtained, one can either remove the redundant models or to reset the zero eigenvalues to a small positive constant.

Under Assumption 2.1, we can invoke the generalized Gauss-Markov-Aitken theorem [16] and find the BLUE by solving the generalized least-square problem for the estimation of the mean, which admits a closed-form solution given by (cf. [17])

$$\hat{\boldsymbol{\mu}}(\mathbf{n}) = (R^T C_\boldsymbol{\varepsilon}^{-1} R)^{-1} R^T C_\boldsymbol{\varepsilon}^{-1} \mathbf{p} = \Psi(\mathbf{n})^{-1} \hat{\boldsymbol{y}}(\mathbf{n}), \quad (4)$$

where $\Psi(\mathbf{n}) \in \mathbb{R}^{\ell \times \ell}$ and $\hat{\boldsymbol{y}}(\mathbf{n}) \in \mathbb{R}^\ell$ are given by

$$\Psi(\mathbf{n}) = R^T C_\boldsymbol{\varepsilon}^{-1} R = \sum_{k \in \mathcal{K}} n_k (R_k)^T (C_k)^{-1} R_k, \quad (5a)$$

$$\hat{\boldsymbol{y}}(\mathbf{n}) = R^T C_\boldsymbol{\varepsilon}^{-1} \mathbf{p} = \sum_{k \in \mathcal{K}} (R_k)^T (C_k)^{-1} \sum_{j \in \mathcal{J}^k} \mathbf{p}^k(\omega_j^k). \quad (5b)$$

Equation (4) is the BLUE for the full vector $\boldsymbol{\mu}$. Since in practice we are only interested in estimating $\mu_1 = \mathbb{E}[p_1]$, we can write the BLUE for μ_1 and its variance as

$$\mathbb{E}[p_1] \approx \mathbf{e}_1^T \hat{\boldsymbol{\mu}}(\mathbf{n}) = \mathbf{e}_1^T \Psi(\mathbf{n})^{-1} \hat{\mathbf{y}}(\mathbf{n}), \quad \mathbb{V}[\mathbf{e}_1^T \hat{\boldsymbol{\mu}}(\mathbf{n})] = \mathbf{e}_1^T \Psi(\mathbf{n})^{-1} \mathbf{e}_1, \quad (6)$$

where $\mathbf{e}_1 = [1, 0, \dots, 0]^T \in \mathbb{R}^\ell$. It can be shown [16, 17] that this estimator for $\mathbb{E}[p_1]$ is also a BLUE and well-defined as long as \mathbf{e}_1 is in the column space of R^T , $\text{col}(R^T)$, which is always satisfied as long as we are sampling the high-fidelity model at least once.

To complete the construction of a MLBLUE, we must find the optimal model selection and the optimal number of samples to draw for each model group. Note that both problems can be addressed by finding the optimal value of \mathbf{n} while allowing some of its entries to be zero: a model that is sampled zero times is this way automatically discarded. To obtain the MOSAP, we minimize the BLUE variance in (6) with respect to \mathbf{n} for a fixed positive computational budget b , i.e., we solve

$$\min_{\mathbf{n} \geq \mathbf{0}} \mathbf{e}_1^T \Psi^{-1}(\mathbf{n}) \mathbf{e}_1, \quad \text{s.t.} \quad \mathbf{n}^T \mathbf{c} \leq b, \quad (7)$$

where $\mathbf{c} \in \mathbb{R}^{|\mathcal{G}|}$ is the cost vector such that c_k is the cost of drawing one realization of model group k . Problem (7) is typically solved by allowing \mathbf{n} to take real values, and then projecting the result to integers. After continuous relaxation, problem (7) becomes a weakly convex nonlinear optimization problem [8]. In fact, it can be shown that even when $\Psi(\mathbf{n})$ is nonsingular for all \mathbf{n} , the Hessian of the objective function in (7) has rank at most ℓ , typically much less than the number of model groups $|\mathcal{G}|$. We remark that for the feasible set of (7) to be non-empty the budget must satisfy

$$b \geq \min_{k: \mathcal{M}_1 \in \mathcal{G}_k} c_k, \quad (8)$$

i.e., we have enough budget to sample the cheapest group containing \mathcal{M}_1 at least once. The resulting algorithm for constructing a MLBLUE is presented in Algorithm 1.

Remark 2.5. We remark that problem (7) is only weakly convex and, as such, admits an infinite set of optimal solutions that are feasible and have the same objective value (cf. [8]). As a consequence, MLBLUEs are not unique.

Algorithm 1 Model and sample allocation optimal MLBLUE

1. Given a set \mathcal{G} of allowed model groupings, estimate the group correlations $\{C_k\}_{k \in \mathcal{K}}$, e.g., by taking $n_{\text{pilot}} \in \mathbb{N}$ pilot samples of all models.
 2. Solve the MLBLUE MOSAP (7) (or the improved MOSAP formulations we present in the next sections) to obtain the optimal sample vector \mathbf{n}^* .
 3. Sample the model groups according to \mathbf{n}^* and stack the samples into a vector \mathbf{p}^* .
 4. Compute the MLBLUE as $\hat{\boldsymbol{\mu}}_1^* = \mathbf{e}_1^T \Psi(\mathbf{n}^*)^{-1} R^T C_\varepsilon^{-1} \mathbf{p}^*$, cf equations (5) and (6).
-

Problem (7) may be affected by severe ill-conditioning. The reason is that $\Psi(\mathbf{n})$ is, in general, only positive semi-definite and becomes singular (see next section) whenever a model is discarded. Schaden and Ullmann address this problem by adding a shift and replacing $\Psi(\mathbf{n})$ by $\Psi_\delta(\mathbf{n}) = \Psi(\mathbf{n}) + \delta I$ for some $\delta > 0$, where the value of δ must be carefully chosen: a large δ artificially increases all entries of \mathbf{n} and leads to oversampling, while a small δ leads to ill-conditioning. (Although the matrices $\Psi_\delta(\mathbf{n})$ are nonsingular for all \mathbf{n} , the rank of the Hessian of the objective function in (7) remains at most ℓ , so the problem is still highly degenerate.) We will show in the next section how we can reformulate MLBLUEs so that no shift parameter is required and so that no issues arise when Ψ is singular or when models are discarded.

Remark 2.6. The number of possible model groupings grows exponentially fast with the number of models ℓ , which might make the MOSAP solution challenging. For instance a model set of size $\ell > 20$ would lead to \mathbf{n} having over a million entries. A solution in this case is to restrict the set \mathcal{G} of allowed groupings. For instance Schaden and Ullmann [8] restrict the maximum size of the allowed groupings to a fixed integer χ and show that in practice little accuracy or efficiency is lost by only using groups of size at most five.

MLBLUEs have three advantages compared with other multilevel/multifidelity estimators. 1) MLBLUEs are provably optimal across all multilevel/multifidelity methods. 2) Unlike most multilevel/multifidelity methods, MLBLUEs are flexible in the sense that they can work with any selection of model groupings and with any model set structure. In contrast, methods such as MLMC and MFMC are restricted to specific model combinations (see Example 2.1). 3) The MLBLUE optimal model selection is automatically found together with the optimal sample allocation without need for a brute-force enumeration approach.

As we will see in Section 4, the multi-output MLBLUEs we introduce inherit the same advantages as the single-output case. We will also see how, thanks to the contributions of Sections 3 and 4, we can add a fourth advantage: 4) The construction of multi-output MLBLUEs and the solution of their MOSAPs can be performed reliably and efficiently. Furthermore, we will numerically demonstrate in Section 5 how MLBLUE is robust to the grouping and sample restrictions arising in the common scenario in which the higher-fidelity models are prohibitive to sample.

Example 2.1. Let us clarify point 2) with an example. We take $\ell = 3$ and, without loss of generality, we order the models by cost. Assuming that the best MLMC and MFMC estimators that can be constructed for this problem use all three models, we obtain the estimators

$$\begin{aligned}\hat{\mu}_{\text{MLMC}} &= \frac{1}{n_1} \sum_{i=1}^{n_1} (p_1(\omega_i^1) - p_2(\omega_i^1)) + \frac{1}{n_2} \sum_{i=1}^{n_2} (p_2(\omega_i^2) - p_3(\omega_i^2)) + \frac{1}{n_3} \sum_{i=1}^{n_3} p_3(\omega_i^3), \\ \hat{\mu}_{\text{MFMC}} &= \frac{1}{m_1} \sum_{i=1}^{m_1} p_1(\omega_i) + \alpha_2 \left(\frac{1}{m_2} \sum_{i=1}^{m_2} p_2(\omega_i) - \frac{1}{m_1} \sum_{i=1}^{m_1} p_2(\omega_i) \right) + \alpha_3 \left(\frac{1}{m_3} \sum_{i=1}^{m_3} p_3(\omega_i) - \frac{1}{m_2} \sum_{i=1}^{m_2} p_3(\omega_i) \right),\end{aligned}$$

where $0 < n_1, n_2, n_3, 0 < m_1 < m_2 < m_3$, and α_2 and α_3 are real parameters that depend on the covariances between the model outputs. We refer the reader to [4] and [3] for further details on MLMC and MFMC respectively. Recalling that samples with different input parameters are independent, it is readily seen from the above estimators that MLMC and MFMC use the following model groupings:

$$\begin{aligned}\mathcal{G}_{\text{MLMC}} &= \{\{\mathcal{M}_1, \mathcal{M}_2\}, \{\mathcal{M}_2, \mathcal{M}_3\}, \{\mathcal{M}_3\}\}, \\ \mathcal{G}_{\text{MFMC}} &= \{\{\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3\}, \{\mathcal{M}_2, \mathcal{M}_3\}, \{\mathcal{M}_3\}\},\end{aligned}$$

i.e., MLMC combines models in pairs, while MFMC samples all models together. We remark that the same holds for larger ℓ : MLMC must work with pairs, and MFMC must be able to sample all models together. If this is not possible, e.g., because estimating correlations is too expensive or, in the multi-output case, simply because a model does not yield that QoI, then the model set must be pruned and adjusted so that an estimator can be constructed. On the other hand, MLBLUEs do not have such restrictions (nor most of the other model restrictions of these methods, cf. [4, 3]): it can work with any set of allowed groupings without needing to discard any models, and it will automatically find the best groups for the estimation problem at hand together with the optimal sample allocation. Generally speaking, while the model set structure could significantly affect the performance of other estimators (see e.g., [7] for a discussion), MLBLUE always retains its optimality. This is why we say that MLBLUEs are flexible.

3. Semidefinite programming reformulation of the MLBLUE model selection and sample allocation problem

While the MLBLUE MOSAP is (weakly) convex, it is still non-trivial to solve in practice, especially when $|\mathcal{G}|$ is large. Near-singularity of $\Psi(\mathbf{n})$ causes the problem to be ill-conditioned, even when the shift

δ is added to form $\Psi_\delta(\mathbf{n})$. Efficient methods should use second-order information, which can be expensive to calculate. In this section, we show how the solution of the MLBLUE MOSAP can be simplified and accelerated by reformulating (7) as a semidefinite programming (SDP) problem.

Semidefinite programming is a powerful paradigm for formulating and solving a wide range of problems in statistics, data science, and control. The problem class most relevant to the MOSAP is experimental design (see for example [18, Section 7.5]) in which the objective is some function of a matrix that is a weighted sum of rank-1 matrices, with the weights being the variables in the problem. (In the MOSAP, the matrices in the sum (5) that defines $\Psi(\mathbf{n})$ are not rank-1, but the SDP formulation techniques still apply.) Efficient interior-point algorithms have been available for SDP since the mid-1990s, and are now included in commercial packages such as MOSEK [13] and open-source packages such as CVXOPT [19].

The first step in the reformulation is to replace Ψ^{-1} with the Moore-Penrose pseudo-inverse Ψ^\dagger so that

$$\hat{\boldsymbol{\mu}} = \Psi^\dagger \hat{\mathbf{y}}, \quad C(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\mu}}) = \Psi^\dagger. \quad (9)$$

This is a standard procedure whenever Ψ is singular and still leads to a BLUE for $\boldsymbol{\mu}$ under Assumption 2.1 [20, 17], but has the advantage of not requiring the δ shift parameter used in [8] (cf. Section 2) without leading to ill-conditioning or ill-posedness of the resulting MOSAP. However, using the pseudo-inverse on its own is still not sufficient to ensure that $\mathbf{e}_1^T \boldsymbol{\mu}$ is also BLUE. As previously said, we also need \mathbf{e}_1 to be in the column space of R^T [17]. This property can be imposed by requiring that model \mathcal{M}_1 is sampled at least once. Problem (7) can then be equivalently reformulated as

$$\min_{\mathbf{n} \geq \mathbf{0}} \mathbf{e}_1^T \Psi^\dagger(\mathbf{n}) \mathbf{e}_1, \quad \text{s.t.} \quad \mathbf{n}^T \mathbf{c} \leq b, \quad \mathbf{n}^T \mathbf{h} \geq 1, \quad (10)$$

where $\mathbf{h} \in \mathbb{R}^{|\mathcal{G}|}$ is a known boolean vector such that $h_k = 1$ if and only if $\mathcal{M}_1 \in \mathcal{G}_k$. The new constraint on the right-hand side imposes that the high-fidelity model is sampled at least once and thus ensures that the estimator is a well-posed BLUE. Since it is necessary for well-posedness of the resulting estimator, this constraint does not exclude any sensible solution of the MOSAP. In particular, the standard Monte Carlo estimator is still feasible for problem (10), ensuring that the feasible set is non-empty (also see Theorem 3.6 in [8]).

Before proceeding we need two auxiliary results. Under Assumption 2.1, Schaden and Ullmann [8] proved that if \mathbf{p} contains samples from all ℓ models, then Ψ is positive definite. We further extend their result to show that Ψ can only be positive definite if all models are sampled, as well as providing a characterization of the null-space of Ψ .

Theorem 3.1. *Let Assumption 2.1 hold. Then Ψ is positive definite if and only if \mathbf{p} contains samples from all ℓ models. More specifically, let $\mathcal{V}_k = \{\mathbf{v} \in \mathbb{R}^\ell : v_i = 0 \ \forall i \text{ s.t. } \mathcal{M}_i \in \mathcal{G}_k\}$, then the null-space of Ψ satisfies $\text{null}(\Psi) = \bigcap_{n_k > 0} \mathcal{V}_k$. In particular, if model $\mathcal{M}_i \in \mathcal{M}$ is never sampled, then the i -th row and column of Ψ and Ψ^\dagger are zero.*

Proof. Let $B_k = (R_k)^T (C_k)^{-1} R_k$ for all k . Since C_k is positive definite for all k , then B_k is positive semidefinite, and so is Ψ , cf. (5). Since $(C_k)^{-1}$ is positive definite and R_k has full row rank, $\mathbf{v} \in \mathbb{R}^\ell$ is in the null-space of B_k if and only if $R_k \mathbf{v} = \mathbf{0}$, which only happens if $\mathbf{v} \in \mathcal{V}_k$. We then conclude that Ψ is singular if and only if there exists $\mathbf{v} \in \mathbb{R}^\ell$ such that (recall: $\mathcal{K} = \{1, \dots, |\mathcal{G}|\}$)

$$\mathbf{v}^T \Psi \mathbf{v} = \sum_{k \in \mathcal{K}} n_k \mathbf{v}^T B_k \mathbf{v} = 0 \quad \Leftrightarrow \quad \forall k \in \mathcal{K} \text{ either } n_k = 0 \text{ or } \mathbf{v} \in \mathcal{V}_k, \quad (11)$$

since n_k and $\mathbf{v}^T B_k \mathbf{v}$ are both non-negative for all k . In particular, this means that the null-space of Ψ is given by $\bigcap_{n_k > 0} \mathcal{V}_k$.

We first prove that if a model is not sampled, then Ψ is singular. Let \mathbf{e}_i with $(\mathbf{e}_i)_j = \delta_{ij}$, (δ_{ij} is the Kronecker delta) be the canonical basis vector. If a model \mathcal{M}_i is not sampled, we have that $n_k = 0$ for all k such that $\mathcal{M}_i \in \mathcal{G}_k$, and $\mathbf{e}_i \in \mathcal{V}_k$ for all k such that $\mathcal{M}_i \notin \mathcal{G}_k$, therefore \mathbf{e}_i is in the nullspace of Ψ , which is therefore singular with zeros in its i -th row and column. The same property holds for Ψ^\dagger , since Ψ and Ψ^\dagger have the same null-space.

We now prove the reverse statement and we show that if all models are sampled, then Ψ is invertible. In this case, for all k corresponding to sampled groups we have $n_k > 0$, and it is readily seen that since we are sampling all models $\bigcap_{n_k > 0} \mathcal{V}_k = \{\mathbf{0}\}$. Hence the null-space of Ψ is trivial and Ψ is invertible. \square

The second auxiliary ingredient is the following result by Albert [21]:

Theorem 3.2 (Result (i) in Theorem 1 in [21]). *Let $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{m \times m}$ be symmetric, let $C \in \mathbb{R}^{n \times m}$, and let*

$$\Phi = \begin{bmatrix} A & C \\ C^T & B \end{bmatrix}. \quad (12)$$

Then $\Phi \succeq 0$ if and only if $A \succeq 0$, $AA^\dagger C = C$ and $B - C^T A^\dagger C \succeq 0$.

Let $t \in \mathbb{R}$ be a scalar slack variable. We are now ready to reformulate (10) as a SDP as follows:

$$\min_{t, \mathbf{n} \geq \mathbf{0}} t, \quad \text{s.t.} \quad \Phi(t, \mathbf{n}) = \begin{bmatrix} \Psi(\mathbf{n}) & \mathbf{e}_1 \\ \mathbf{e}_1^T & t \end{bmatrix} \succeq 0, \quad \mathbf{n}^T \mathbf{c} \leq b, \quad \mathbf{n}^T \mathbf{h} \geq 1. \quad (13)$$

This is indeed an SDP since $\Psi(\mathbf{n})$ is linear in \mathbf{n} and thus $\Phi(t, \mathbf{n})$ is linear in (t, \mathbf{n}) .

Theorem 3.3. *Let Assumption 2.1 hold. Then problems (10) and (13) are equivalent.*

Proof. By Theorem 3.2 we have that

$$\Phi \succeq 0 \quad \Leftrightarrow \quad \Psi \succeq 0, \quad \Psi \Psi^\dagger \mathbf{e}_1 = \mathbf{e}_1, \quad t \geq \mathbf{e}_1^T \Psi^\dagger \mathbf{e}_1. \quad (14)$$

The first of the three conditions is automatically satisfied since Ψ is non-negative definite. To complete the proof it is sufficient to prove that the second condition also always holds. In fact, this would imply that $\Phi \succeq 0$ if and only if $t \geq \mathbf{e}_1^T \Psi^\dagger \mathbf{e}_1$, and taking the minimum of the latter expression over t would then yield the thesis, since $\min_t t = \mathbf{e}_1^T \Psi^\dagger \mathbf{e}_1$. Note that the spaces $\text{col}(\Psi)$ and $\text{row}(\Psi)$ coincide since Ψ is symmetric. For the second condition on the right-hand side of (14) to hold it is sufficient that $\mathbf{e}_1 \in \text{col}(\Psi)$ since $\Psi \Psi^\dagger$ is the orthogonal projector onto $\text{col}(\Psi)$. Note that since $\text{col}(\Psi) = \text{row}(\Psi) = \text{null}(\Psi)^\perp$, it is sufficient to demonstrate that \mathbf{e}_1 is orthogonal to $\text{null}(\Psi)$, which, by Theorem 3.1, is given by $\text{null}(\Psi) = \bigcap_{n_k > 0} \mathcal{V}_k$. We observe that the last constraint in (10) and (13) ensures that model \mathcal{M}_1 is sampled at least once. Hence, if we let \bar{k} index the sampled groups (i.e., those for which $n_{\bar{k}} > 0$) that include \mathcal{M}_1 we have

$$\text{null}(\Psi) = \bigcap_{n_k > 0} \mathcal{V}_k \subseteq \bigcap_{\bar{k}} \mathcal{V}_{\bar{k}} \subseteq \{\mathbf{v} \in \mathbb{R}^\ell : v_1 = 0\} = \text{span}(\mathbf{e}_1)^\perp. \quad (15)$$

Hence $\mathbf{e}_1 \in \text{null}(\Psi)^\perp = \text{row}(\Psi) = \text{col}(\Psi)$, and hence $\Psi \Psi^\dagger \mathbf{e}_1 = \mathbf{e}_1$ and the theorem is proved. \square

We have now proved that the MLBLUE MOSAP can be formulated as the SDP (13), which can be solved reliably and efficiently. In practice it is also often useful to solve an alternative MOSAP in which the user prescribes a statistical error tolerance ε^2 and minimizes the computational expense subject to this restriction. This new MOSAP reads

$$\min_{\mathbf{n} \geq \mathbf{0}} \mathbf{n}^T \mathbf{c}, \quad \text{s.t.} \quad \mathbf{e}_1^T \Psi^\dagger(\mathbf{n}) \mathbf{e}_1 \leq \varepsilon^2, \quad \mathbf{n}^T \mathbf{h} \geq 1, \quad (16)$$

which is a straightforward modification of (10). Problem (16) appears harder to solve than (10) since the nonlinear variance function now appears as a constraint. Again, we can reformulate (16) as a SDP as follows

$$\min_{\mathbf{n} \geq \mathbf{0}} \mathbf{n}^T \mathbf{c}, \quad \text{s.t.} \quad \Phi(\varepsilon^2, \mathbf{n}) = \begin{bmatrix} \Psi(\mathbf{n}) & \mathbf{e}_1 \\ \mathbf{e}_1^T & \varepsilon^2 \end{bmatrix} \succeq 0, \quad \mathbf{n}^T \mathbf{h} \geq 1. \quad (17)$$

Corollary 3.4. *Let Assumption 2.1 hold. Then problems (16) and (17) are equivalent.*

Proof. In the proof of Theorem 3.3 we have already shown that $\Phi(t, \mathbf{n}) \succeq 0$ if and only if $\mathbf{e}_1^T \Psi^\dagger(\mathbf{n}) \mathbf{e}_1 \leq t$. Setting $t = \varepsilon^2$ yields the thesis. \square

Note that (17) is also a SDP and it is as easy to solve as (13). In fact, the two problems are strongly related: by introducing a parameter $\tau > 0$ we can trace out the Pareto frontier defined by the two objectives of statistical error and cost as follows:

$$\min_{t, \mathbf{n} \geq \mathbf{0}} t + \tau \mathbf{n}^T \mathbf{c}, \quad \text{s.t.} \quad \Phi(t, \mathbf{n}) = \begin{bmatrix} \Psi(\mathbf{n}) & \mathbf{e}_1 \\ \mathbf{e}_1^T & t \end{bmatrix} \succeq 0, \quad \mathbf{n}^T \mathbf{h} \geq 1. \quad (18)$$

There are values of τ that produce the solutions of (13) and (17). This problem (18) is also a SDP and can be solved as easily as the other two formulations.

Remark 3.1. It was shown in [8] that some approximate control variate methods, such as the ACV-IS estimator of [7], are BLUEs for a specific set of allowed model groupings. However, it is nontrivial to solve the sample allocation problem in [7] numerically. Formulations (13) and (17) can be used to accelerate and simplify the set up of *any* estimator that is also a MLBLUE.

Remark 3.2. The SDP formulations of MOSAP in (13), (17), and (18) can incorporate additional linear constraints (e.g., on the number of samples) while remaining SDPs. For instance, in Section 5 we consider a scenario in which the high-fidelity models are prohibitive to sample and the maximum number of high-fidelity samples is thus restricted.

An advantage of formulating the MOSAPs as SDPs is that they can easily be extended to the multiple output case without impacting their numerical solvability, as we show next.

4. Multi-output multilevel best linear unbiased estimators

In this section we extend MLBLUEs to the multi-output case and introduce SDP formulations for their MOSAPs. We tackle the multi-output estimation problem by simultaneously constructing m MLBLUEs, one for each output, and setting up two multiple-output SDP MOSAPs, the first minimizing the total cost, and the second minimizing the maximum statistical error.

Consider a single model set $\mathcal{M} = \{\mathcal{M}_i\}_{i=1}^\ell$, corresponding now to m different output sets indexed by $s = 1, \dots, m$, so that p_i^s is the output of the i -th model corresponding to the s -th QoI. We do not assume that all the outputs p_i^s can actually be sampled, since models might not necessarily produce all outputs. For this reason, we construct m separate model coupling sets $\mathcal{G}^s \subseteq 2^{\mathcal{M}} \setminus \emptyset$ for $s = 1, \dots, m$, each containing the possible or feasible model groupings for the s -th QoI.

We now simultaneously construct m MLBLUEs, each starting from its feasible set \mathcal{G}^s obtaining the estimators

$$\mathbb{E}[p_1^s] \approx \mathbf{e}_1^T \hat{\boldsymbol{\mu}}^s = \mathbf{e}_1^T \Psi_s^\dagger \hat{\mathbf{y}}_s, \quad \mathbb{V}[p_1^s] = \mathbf{e}_1^T \Psi_s^\dagger \mathbf{e}_1, \quad s = 1, \dots, m, \quad (19)$$

where $\hat{\mathbf{y}}_s, \Psi_s$ are constructed in the same way as in the single-output case. We still consider each Ψ_s to be of size ℓ -by- ℓ even if not all ℓ models are used: if model i cannot be used to estimate the s -th QoI, we simply set the i -th row and column of Ψ_s to zero. With this strategy, we can incorporate models that output only some of the QoIs within the standard MLBLUE construction, so that optimality is still ensured. Note that a Ψ_s constructed this way remains consistent with the properties of the matrix Ψ discussed in Theorem 3.1.

To set up all m estimators concurrently, we need to design and solve a multi-output MOSAP. We have three options: 1) fix a computational budget and minimize the maximum estimator variance. 2) Fix a statistical error tolerance ε_s^2 for each output and minimize the total cost. 3) Trace out a Pareto frontier defined by the two objectives of maximum estimator variance and cost. The resulting formulations are as follows.

1) *Minimizing the maximum estimator variance for a fixed budget.* We write the multi-output MOSAP as

$$\min_{t, \mathbf{n} \geq \mathbf{0}} t, \quad \text{s.t.} \quad \mathbf{e}_1^T \Psi_s^\dagger(\mathbf{n}) \mathbf{e}_1 \leq t, \quad \mathbf{n}^T \mathbf{h}^s \geq 1, \quad \mathbf{n}^T \mathbf{c} \leq b, \quad s = 1, \dots, m. \quad (20)$$

Its corresponding SDP reformulation reads:

$$\min_{t, \mathbf{n} \geq \mathbf{0}} t, \quad \text{s.t.} \quad \Phi_s(t, \mathbf{n}) = \begin{bmatrix} \Psi_s(\mathbf{n}) & \mathbf{e}_1 \\ \mathbf{e}_1^T & t \end{bmatrix} \succeq 0, \quad \mathbf{n}^T \mathbf{h}^s \geq 1, \quad \mathbf{n}^T \mathbf{c} \leq b, \quad s = 1, \dots, m. \quad (21)$$

We now have m different positive semidefiniteness constraints, one for each output, but otherwise this is similar to the single-output MOSAPs (10) and (13). Let $\mathcal{G} = \bigcup_{s=1}^m \mathcal{G}^s$ be the set of all possible model groupings across all estimators. Then the vector \mathbf{n} is of size $|\mathcal{G}|$ and the vectors $\mathbf{h}^s \in \mathbb{R}^{|\mathcal{G}|}$ are constructed in the same way as before: for all $\mathcal{G}_k \in \mathcal{G}$, set $\mathbf{h}_k = 1$ if $\mathcal{M}_1 \in \mathcal{G}_k \in \mathcal{G}^s$ and $\mathbf{h}_k = 0$ otherwise. Following the same approach as in the proof of Theorem 3.3 it is easy to see that for any pair (t, \mathbf{n}) that is feasible for problem (21) we have that $\mathbf{e}_1^T \Psi_s^\dagger(\mathbf{n}) \mathbf{e}_1 \leq t$ for all s , and therefore the two optimization problems (20) and (21) are equivalent.

Remark 4.1. Minimizing the maximum statistical error across all outputs as in (21) is not the only option and other objective functions can be considered. For instance, if we introduce a slack vector $\mathbf{t} \in \mathbb{R}^m$, with t_s corresponding to the statistical error of output s , we can replace $\Phi_s(t, \mathbf{n})$ in (21) with $\Phi_s(\mathbf{t}_s, \mathbf{n})$, and then minimize a linear or convex quadratic function³ of \mathbf{t} , e.g., $\|\mathbf{t}\|_1$, $\|\mathbf{t}\|_{\mathbf{w}} = \sum w_s t_s$ for some positive vector of weights \mathbf{w} , or $\|\mathbf{t}\|_2$.

2) *Minimizing the total cost for fixed statistical error tolerances.* An alternative MOSAP formulation can be obtained by setting a separate statistical error tolerance $\varepsilon_s^2 > 0$ for each output s , and then minimizing the total cost of all m estimators. This yields the two equivalent formulations:

$$\min_{\mathbf{n} \geq \mathbf{0}} \mathbf{n}^T \mathbf{c}, \quad \text{s.t.} \quad \mathbf{e}_1^T \Psi_s^\dagger(\mathbf{n}) \mathbf{e}_1 \leq \varepsilon_s^2, \quad \mathbf{n}^T \mathbf{h}^s \geq 1, \quad s = 1, \dots, m, \quad (22)$$

$$\min_{\mathbf{n} \geq \mathbf{0}} \mathbf{n}^T \mathbf{c}, \quad \text{s.t.} \quad \Phi_s(\varepsilon_s^2, \mathbf{n}) = \begin{bmatrix} \Psi_s(\mathbf{n}) & \mathbf{e}_1 \\ \mathbf{e}_1^T & \varepsilon_s^2 \end{bmatrix} \succeq 0, \quad \mathbf{n}^T \mathbf{h}^s \geq 1, \quad s = 1, \dots, m. \quad (23)$$

The formulations (22) and (23) extend (16) and (17), respectively, to the multi-output case.

3) *Trace out the Pareto frontier between the two objectives of maximum estimator variance and total cost.* A third possibility is to introduce a nonnegative parameter τ and consider the MOSAP

$$\min_{t, \mathbf{n} \geq \mathbf{0}} t + \tau \mathbf{n}^T \mathbf{c}, \quad \text{s.t.} \quad \mathbf{e}_1^T \Psi_s^\dagger(\mathbf{n}) \mathbf{e}_1 \leq t, \quad \mathbf{n}^T \mathbf{h}^s \geq 1, \quad s = 1, \dots, m. \quad (24)$$

Its corresponding SDP reformulation reads:

$$\min_{t, \mathbf{n} \geq \mathbf{0}} t + \tau \mathbf{n}^T \mathbf{c}, \quad \text{s.t.} \quad \Phi_s(t, \mathbf{n}) = \begin{bmatrix} \Psi_s(\mathbf{n}) & \mathbf{e}_1 \\ \mathbf{e}_1^T & t \end{bmatrix} \succeq 0, \quad \mathbf{n}^T \mathbf{h}^s \geq 1, \quad s = 1, \dots, m. \quad (25)$$

Equation (25) is the multi-output equivalent of (18). Its solution will match that of (21) for some non-negative value of τ . For the variant of (23) in which we define ε_s to have the same positive value for all $s = 1, 2, \dots, m$, there is also a nonnegative value of τ for which the solutions of (25) and (23) coincide.

We highlight the fact that these multi-output MBLUEs leave the model selection completely to the MOSAP *without need for the user to make any prior decision about the quality of the models or the structure*

³A convex quadratic objective would add a second-order cone constraint to (21); the resulting problem is still a conic program and is no more difficult to solve than the original SDP [22]. For example, to minimize $\|\mathbf{t}\|_2$, we introduce a scalar variable \bar{t} and the constraint $\bar{t} \geq \|\mathbf{t}\|_2$ (which can be expressed as a second-order cone constraint), and define the objective to be \bar{t} .

of the model set. This is a convenient feature which becomes essential when models are heterogeneous and there is no clear distinction as to which ones are more suitable for the joint estimation problem. We also remark that multi-output MLBLUEs are the first instance of a multi-output multilevel or multifidelity estimator that simultaneously constructs a separate estimator for each QoI.

Remark 4.2. Note that the MLBLUEs in (19) are all BLUE if considered *independently*, but they are not BLUE if the multi-output estimation is considered *jointly*. One could account for the correlations between different outputs and construct a single multi-output MLBLUE in which different outputs are treated as different “models”. However, this would lead to a total of up to $m\ell$ possible models, one for each (model, output) combination, causing the number of possible groupings to scale like $O(2^{m\ell})$ and considerably increase the cost of solving the MOSAP.

Remark 4.3. We remark that single and multi-output MLBLUEs (and all multifidelity methods based on covariances) can be extended to handle infinite-dimensional QoI with minor modifications. Provided that the output belongs to a Bochner space $\mathcal{L}^2(\Omega, \mathcal{V})$ for a suitable sample space Ω and a Hilbert space \mathcal{V} , it is sufficient to define the covariance operator (and consequently variance and correlation operators) as the inner product $(\cdot, \cdot)_{\mathcal{L}^2(\Omega, \mathcal{V})}$ of $\mathcal{L}^2(\Omega, \mathcal{V})$, i.e.,

$$C(X, Y) = (X, Y)_{\mathcal{L}^2(\Omega, \mathcal{V})} = \mathbb{E}[(X - \mathbb{E}[X], Y - \mathbb{E}[Y])_{\mathcal{V}}], \quad \forall X, Y \in \mathcal{L}^2(\Omega, \mathcal{V}), \quad (26)$$

where $(\cdot, \cdot)_{\mathcal{V}}$ is the inner product in \mathcal{V} . The above is a natural generalization: by taking $\mathcal{V} = \mathbb{R}$ we recover the standard covariance operator between scalar random variables. This approach has the advantage of yielding scalar-valued covariances that can be used in multifidelity methods in the same way as in the scalar case.

5. Numerical results

In this section we test both the new SDP formulations of the MLBLUE MOSAP and the single- and multi-output MLBLUEs. The aim is to demonstrate the advantages of the new MOSAP formulations and of MLBLUEs themselves in problems that reflect the complexity of typical computational engineering applications. Specifically, we will show how the improved efficiency of MLBLUEs results from their ability to work with any selection of model groupings. Here we only compare MLBLUE with MLMC and MFMC: performance comparisons between MLBLUE and other estimators such as the ACV methods for a single QoI are available in [8, 9] and extending other estimators to multi-output problems is beyond the scope of this work.

We implemented the single- and multi-output MLBLUEs (as well as single- and multi-output MLMC and MFMC) into an open-source Python software, BLUEST, which is publicly available on GitHub at <https://github.com/croci/bluest>. BLUEST allows users to easily wrap Python models (and, by extension, any model that can be called from Python), and supports MPI sample parallelization of serial and MPI-parallelized models.

5.1. Model heterogeneity

We start by considering two problems with heterogeneous model sets: Problem 1 uses global and local grid refinement to define the models while Problem 2 uses global grid and timestep refinement as well as a combination of different PDE and ODE models. The objective here is to demonstrate the improved efficiency of multi-output MLBLUEs and of the new SDP formulations. We first investigate the computational efficiency of the MLBLUE, MLMC, and MFMC methods constructed for each problem and then describe the computational advantages of the new MOSAP formulations.

5.1.1. Problem 1: Steady Navier-Stokes flow past two cylinders

Problem description. As a first test problem we consider a steady flow past two cylindrical obstacles of different sizes, as described by the steady Navier-Stokes equations:

$$\begin{cases} -\nu(\omega)\Delta\mathbf{u} + \mathbf{u}\nabla\mathbf{u} + \nabla q = 0, & \nabla \cdot \mathbf{u} = 0, & \mathbf{x} \in \mathcal{D}_2, \\ \mathbf{u}|_{\Gamma_t} = \mathbf{u}|_{\Gamma_b} = \mathbf{u}|_{\partial\mathcal{C}_1} = \mathbf{u}|_{\partial\mathcal{C}_2} = \mathbf{0}, & \nu(\omega)\nabla\mathbf{u}|_{\Gamma_r}\boldsymbol{\eta} - q|_{\Gamma_r}\boldsymbol{\eta} = 0, \\ \mathbf{u}|_{\Gamma_l} = (4|\Gamma_l|^{-2}u_{\text{in}}(\omega)x_2(|\Gamma_l| - x_2), 0)^T, & \mathbf{x} = (x_1, x_2)^T, \end{cases} \quad (27)$$

where the domain \mathcal{D}_2 is described Figure 1. Here \mathbf{u} is the flow velocity and q its pressure, and $\boldsymbol{\eta}$ denotes the outer unit normal vector of the boundary. The viscosity ν and the maximum inlet velocity u_{in} are modeled as uniform random variables, with $\nu(\omega) \in [5 \times 10^{-4}, 1.5 \times 10^{-3}]$ and $u_{\text{in}} \in [0.1, 0.5]$. This leads to a Reynolds number which is also random (albeit not uniform), and given by

$$\text{Re}(\omega) = \frac{u_{\text{mean}}(\omega) |\mathcal{C}_2|}{\nu(\omega) \pi}, \quad \text{Re}(\omega) \in [7.\bar{1}, 106.\bar{6}].$$

where $u_{\text{mean}} = \frac{2}{3}u_{\text{in}}$ is the mean inlet velocity and $|\mathcal{C}_2|/\pi$ is the diameter of the largest obstacle.




Figure 1: The domain of Problem 1. The rectangle has length $|\Gamma_b| = |\Gamma_t| = 2.2$ and width $|\Gamma_l| = |\Gamma_r| = 0.41$. The two obstacles \mathcal{C}_1 and \mathcal{C}_2 have centers at $(0.2, 0.2)$ and $(1.0, 0.2)$ and radii 0.05 and 0.08 respectively. Γ_l and Γ_r are the inflow and outflow boundaries respectively and the other boundaries are walls.

Quantities of interest. We consider six QoIs: the drag and lift coefficients around each obstacle

$$(c_{\text{drag}}^i, c_{\text{lift}}^i)^T = \frac{2\pi}{u_{\text{mean}}^2 |\mathcal{C}_i|} \int_{\partial\mathcal{C}_i} (\nu\nabla\mathbf{u} - pI)\boldsymbol{\eta} \, ds, \quad i = 1, 2, \quad (28)$$

as well as the difference in pressure before and after each cylinder

$$\Delta_{\mathcal{C}_1} q = q(0.15, 0.2) - q(0.25, 0.2), \quad \Delta_{\mathcal{C}_2} q = q(0.72, 0.2) - q(0.88, 0.2). \quad (29)$$

Model set. We construct twelve grid-based models for this problem as follows. First, we construct three base meshes, one coarse, one medium, and one fine. Then we locally refine each base mesh in three different ways: around both obstacles, around \mathcal{C}_1 only, or around \mathcal{C}_2 only. We then have three base grids with four configurations each (three types of local refinements + no refinement), for a total of twelve grids, which we show in Figure 2. We build the meshes with the open-source software **Gmsh**⁴ (version 4.10.5) and the mesh-generating scripts are available in the **BLUEST** open-source repository⁵. Note that the grids refined around a single obstacle improve the approximation of the QoIs defined on that obstacle, but not of those defined on the other cylinder. We design the model set in this way to investigate the performance of multi-output MLBLUEs in a setting in which some low-fidelity models yield good approximations of only some QoIs. While Problem 1 is a simple example, a conceptually similar setting often arises in multi-physics problems where low-fidelity models are defined by simplifying some of the physics. In such problems, the low-fidelity models will yield poor approximations of the QoIs related to the physics that were simplified.

⁴Available at <https://gmsh.info/>.

⁵See `mesh_generator.py` in the `bluest/examples/navier-stokes` folder in the **BLUEST** repository.




Figure 2: The grids used to define the twelve models for Problem 1. Each column corresponds to a different “base” mesh: fine (left), medium (center), coarse (right). The actual “base” meshes are in the last row.

Covariance estimation. We take 100 pilot samples of all models to estimate model covariances. This is a larger amount than needed in practice in most cases [3]; we choose a high number to sanitize our results from pilot sampling errors.

Solver and model costs. We solve (27) over all grids with the FE method using the Taylor-Hood FE pair and the FEniCS library [23] using a direct linear solver. As to costs, we use the number of degrees of freedom as a proxy for simplicity: we set these pseudo-costs to be $(n_{\text{dofs}}^i)^2$, where n_{dofs}^i is the number of degrees of freedom of model i , and we have assumed a quadratic cost complexity for the nonlinear solver used to solve (27). Going from top to bottom and from left to right in Figure 2, the number of degrees of freedom for each grid are $\{68619, 47952, 53817, 33114\}$ (fine grids), $\{18512, 12875, 14471, 8924\}$ (medium grids) $\{5225, 3932, 4274, 3026\}$ (coarse grids). We do not account for pilot samples in our cost computations since all the methods we compare (MLBLUE, MLMC and MFMC) benefit from more accurate estimates.

Estimator setup. We solve the statistical-error-constrained SDP MOSAP (23) with tolerances $\varepsilon_s = 10^{-3} \sqrt{\mathbb{V}[p_1^s]}$ for $s = 1, \dots, m = 6$, where $\mathbb{V}[p_1^s]$ is estimated from the pilot samples. Here we consider all model groupings of size at most $\chi = 7$ (cf. Remark 2.6). For constructing the MLMC and MFMC estimators, we proceed as follows. For each possible model combination that includes the high-fidelity model, we loop over all QoIs $s = 1, \dots, m$. For each QoI, we order the models (by cost for MLMC [4] and by correlation for MFMC [3]) and compute the QoI optimal sample allocation satisfying the tolerance ε_s by using the MLMC/MFMC single-output analytic expression. We then take the maximum number of samples across all QoIs and use the result to compute the corresponding total multi-output estimator cost for that model combination. Finally, we select the model group corresponding to the minimum total cost. Note that by setting up MLMC and MFMC this way there is no need to provide any model ordering *a priori*. For MFMC, if for any model combination there is a QoI for which the MFMC model ordering constraints [3] are not satisfied, then that model combination is discarded. (Note that MLMC does not have an equivalent restriction on the variances⁶ [4].) In all methods, once a real-valued sample allocation \mathbf{n}^* is found, we project it to an integer-valued solution (cf. Remark 5.2).

Results: total estimation cost. Once the optimal MLBLUE, MLMC and MFMC estimators have been constructed, we compute their total cost so that we can compare their efficiency. Results are shown in Figure 3. MLBLUE is roughly three times more efficient than MLMC and two times more efficient than MFMC for the same statistical error tolerances. The higher efficiency of MLBLUE derives from a better management and exploitation of the cheaper local refinement models: MLBLUE exploits all coarse grid models and only

⁶MLMC variances are often assumed as decaying in MLMC convergence theory. However, this is not a requirement of MLMC itself [4].




Figure 3: Total cost of the MOSAP-optimal MLBLUE, MLMC, and MFMC methods for Problem 1. Colors denote the “base” meshes (darker = finer), while the patterns indicate the local refinements around obstacles.




Figure 4: Points on the Pareto frontier MLBLUE for Problem 1. Here $\tilde{\tau} = \tau \|c\|_2$, where τ is the Pareto parameter from (25) and c is the model cost vector. From right to left, the points in the figure correspond to $\tilde{\tau} = 10^i$ for $i = -7, \dots, 4$.

takes a few samples of the expensive medium grid models. In contrast, MLMC and MFMC draw many medium grid model samples, which account for a considerable portion of their total cost.

Results: Pareto frontier optimization. We solve the Pareto formulation of the MLBLUE MOSAP (25) to analyze the trade-off between the two constraints of budget and accuracy. Figure 4 shows the Pareto-optimal points corresponding to different values of the Pareto parameter τ . We take $\tau = \tilde{\tau} \|\mathbf{c}\|_2^{-1}$, where $\tilde{\tau} = 10^i$ for $i = -7, \dots, 4$. The points are plotted in the total MLBLUE cost vs normalized error plane, where

$$\text{normalized error} = \max_s \left(\frac{\mathbb{V}[\mathbf{e}_1^T \hat{\boldsymbol{\mu}}^s]}{\mathbb{V}[p_1^s]} \right)^{1/2}. \quad (30)$$

For $s = 1, \dots, m$, the high-fidelity model output variances $\mathbb{V}[p_1^s]$ are estimated from pilot samples and $\mathbb{V}[\mathbf{e}_1^T \hat{\boldsymbol{\mu}}^s]$ is the estimator variance computed from the MOSAP solution using (19). For large enough τ , accuracy becomes unimportant and the MOSAP yields the cheapest, least accurate estimator satisfying the constraints in (25), namely, a single Monte Carlo sample of the high-fidelity model. On the other hand, as $\tau \rightarrow 0$, the Pareto frontier becomes a line in the loglog plot with a slope of $-\frac{1}{2}$ as expected for a Monte Carlo simulation: $\text{error} = \text{cost}^{-1/2}$. For the values of τ between these two extremes, the behavior of the frontier is less obvious, and in this low-to-moderate-budget regime, it provides precise quantitative information about the available tradeoffs. Here we have aggregated the MLBLUE error across all QoIs for simplicity. For more precise calibration, it may be beneficial to compute separate Pareto frontiers for each output.

5.1.2. Problem 2: Hodgkin-Huxley model

Problem description - Hodgkin-Huxley high-fidelity model. As a second test problem we consider the Hodgkin-Huxley equations for neuron membrane action potential in 1D on the unit interval:

$$\begin{cases} c(\omega)v_t &= \iota(\omega) + \epsilon(\omega)\Delta v + g_K\alpha^4(v_K - v) + g_{\text{Na}}\beta^3\gamma(v_{\text{Na}} - v) + g_l(v_l - v), & t \in [0, 20]. \\ \alpha_t &= \phi_\alpha(v)(1 - \alpha) - \psi_\alpha(v)\alpha, \\ \beta_t &= \phi_\beta(v)(1 - \beta) - \psi_\beta(v)\beta, \\ \gamma_t &= \phi_\gamma(v)(1 - \gamma) - \psi_\gamma(v)\gamma, \end{cases} \quad (31)$$

$$\begin{aligned} \phi_\alpha(v) &= \frac{1}{10} \frac{1 - v/10}{e^{1-v/10} - 1}, & \psi_\alpha(v) &= \frac{e^{-v/80}}{8}, \\ \phi_\beta(v) &= \frac{2.5 - v/10}{e^{2.5-v/10} - 1}, & \psi_\beta(v) &= 4e^{-v/18}, \\ \phi_\gamma(v) &= 0.07e^{-v/20}, & \psi_\gamma(v) &= (1 + e^{3-v/10})^{-1}. \end{aligned} \quad (32)$$

The boundary conditions are $v_x(t, 0, \omega) = v(t, 1, \omega) = 0$ for all ω, t (homogeneous Neumann and Dirichlet on the left and right boundary respectively), and the initial conditions are given by

$$v(0, x, \omega) = v_{\text{eq}}, \quad w(0, x, \omega) = \frac{\phi_w(v - v_{\text{eq}})}{\phi_w(v - v_{\text{eq}}) + \psi_w(v - v_{\text{eq}})}, \quad \text{for } w \in \{\alpha, \beta, \gamma\}, \quad (33)$$

where v_{eq} is the equilibrium potential, whose value is given in Table 1 together with the values of all the deterministic parameters used for this problem. We refer the reader to [24] for a derivation and description of the original Hodgkin-Huxley model in ordinary differential equation (ODE) form and its PDE extension, which was studied in [25].

In (31), $v = v(t, x, \omega)$ is the action potential across a neuron membrane, which depends on the difference in the concentration of sodium (Na) and potassium (K) ions inside and outside of the neuron cell. These differences in concentrations are maintained by active transport mechanisms such as sodium-potassium pumps. In (31), these mechanisms are modeled by the variables $\alpha = \alpha(t, x, \omega)$, $\beta = \beta(t, x, \omega)$, $\gamma = \gamma(t, x, \omega)$, with $0 \leq \alpha, \beta, \gamma \leq 1$, which represent the activation and inactivation of gated ion transport channels. The variable α is the potassium gated channel activation, and the variables β and γ are the sodium gated channel

Parameter	Meaning	Value
g_{Na}	Sodium conductance p.u.a.	120 mScm ⁻²
g_{K}	Potassium conductance p.u.a.	36 mScm ⁻²
g_l	Leakage conductance p.u.a.	0.3 mScm ⁻²
v_{Na}	Sodium Nernst potential	56 mV
v_{K}	Potassium Nernst potential	-77 mV
v_l	Leakage Nernst potential	-60 mV
v_{eq}	Equilibrium potential	-67.38614 mV
c_0	Membrane capacitance p.u.a.	1 μFcm^{-2}
ϵ_0	Neuron fiber conductance	0.33616 mS
ι_0	Applied current p.u.a.	33.6 mAcm ⁻²

Table 1: *Parameters used in the Hodgkin-Huxley model (31) and their values. The units S, V, F, and A indicate siemens, volts, farads and amperes respectively. The acronym “p.u.a.” stands for “per unit area”.*

activation and inactivation respectively. Additionally, we model the membrane capacitance $c(\omega)$, the neuron fiber conductance $\epsilon(\omega)$ and the applied current per unit area $\iota(\omega)$ as random variables defined by:

$$c(\omega) = (c_0 + 0.2z_1^2(\omega)) \mu\text{Fcm}^{-2}, \quad \epsilon(\omega) = e^{z_2} \text{ mS}, \quad \iota(\omega) = \iota_0 \times 5^{2v(\omega)-1} \text{ mAcm}^{-2}, \quad (34)$$

where $z_1(\omega)$ and $v(\omega)$ are standard Gaussian and uniform random variables respectively, and $z_2(\omega)$ is a Gaussian random variable with mean and variance chosen so that $\mathbb{E}[e^{z_2}] = \epsilon_0$ and $\mathbb{V}[e^{z_2}]^{1/2} = 0.2\epsilon_0$.

Problem description - FitzHugh-Nagumo low-fidelity model. As a low-fidelity model, we take the FitzHugh-Nagumo model, which is obtained by introducing the simplifying assumptions that β and $\alpha + \gamma$ are constant in time and space with $\beta = \bar{\beta}$ and $\alpha + \gamma = \bar{\gamma}$. The values of $\bar{\beta}$ and $\bar{\gamma}$ are obtained from the initial conditions (33). The resulting FitzHugh-Nagumo equations read

$$\begin{cases} c(\omega)v_t &= \iota(\omega) + \epsilon(\omega)\Delta v + g_{\text{K}}\alpha^4(v_{\text{K}} - v) + g_{\text{Na}}\bar{\beta}^3(\bar{\gamma} - \alpha)(v_{\text{Na}} - v) + g_l(v_l - v), \\ \alpha_t &= \phi_\alpha(v)(1 - \alpha) - \psi_\alpha(v)\alpha, \end{cases} \quad (35)$$

where the initial and boundary conditions for v and α are the same as for the Hodgkin-Huxley model. We refer the reader to [24] for further information about this model.

Quantities of interest. As QoIs, we choose the peak membrane potential $v_{\text{peak}}(\omega) = \max_{t,x} v(t, x, \omega)$, and the net total membrane, potassium, sodium, and leakage currents, which are respectively given by:

$$\iota_{\text{me}}(\omega) = \int_0^{20} \int_0^1 (\iota(\omega) + \epsilon(\omega)\Delta v) \, dx \, dt, \quad \iota_{\text{Na}}(\omega) = \int_0^{20} \int_0^1 g_{\text{Na}}\bar{\beta}^3 d(v_{\text{Na}} - v) \, dx \, dt, \quad (36)$$

$$\iota_{\text{K}}(\omega) = \int_0^{20} \int_0^1 g_{\text{K}}\alpha^4(v_{\text{K}} - v) \, dx \, dt, \quad \iota_l(\omega) = \int_0^{20} \int_0^1 g_l(v_l - v) \, dx \, dt. \quad (37)$$

Model set. We consider four base models: the Hodgkin-Huxley PDE model (31), the FitzHugh-Nagumo PDE model (35), and the corresponding ODE models obtained by setting $\epsilon = 0$ and taking v , α , β , and γ to be constant in space. From each of these models, we consider three different types of grid and timestep refinements: we discretize the equations in time and space (PDE models only) using uniform grids with mesh size and timestep respectively given by $\Delta x = 2^{-(2i+1)}$ and $\Delta t = \frac{2}{5}\Delta x$ for $i = 1, 2, 3$. This yields a model set comprised of twelve models in total, in which the Hodgkin-Huxley PDE model with the finest grid and timestep is the high-fidelity model. This model set is an example of a heterogeneous model set in which again not all models are good for all outputs. For instance, the ODE models are unlikely to accurately approximate ι_{me} (the diffusion term is zero for the ODEs), and the FitzHugh-Nagumo models are unlikely to yield accurate samples of ι_{Na} (it involves quantities approximated as constants).

Covariance estimation. As for Problem 1, we take 100 pilot samples of all models to estimate model correlations.

Solver and model costs. We discretize the PDEs in space with the FE method using piecewise linear elements in **FEniCS** [23] and both ODEs and PDEs in time using the backward Euler method. We set the model costs to be given by $\tilde{c}_i n_{\text{ts}}^i n_{\text{dofs}}^i$, where n_{ts}^i and n_{dofs}^i are the number of timesteps and of degrees of freedom used by model i respectively (n_{dofs}^i is the number of equations in the ODE models). Here the \tilde{c}_i factor is 1 for the ODE models and 8 for the PDE models, and accounts for the fact that PDEs have a slightly higher cost per degree-of-freedom than ODEs. The values of \tilde{c}_i were estimated from CPU times.

Estimator setup. We use the same setup as for Problem 1. Again, we set the statistical error tolerances for the SDP MOSAP to be $\varepsilon_s = 10^{-3} \sqrt{\mathbb{V}[p_1^s]}$ for $s = 1, \dots, m = 4$ where $\mathbb{V}[p_1^s]$ is estimated from pilot samples.




Figure 5: Total cost of the MOSAP-optimal MLBLUE, MLMC, and MFMC methods for Problem 2. The three darkest colors correspond to the PDE models with darker tints indicating finer grids and timesteps. The lightest color corresponds to the total computational effort spent on all the ODE models. Here we do not distinguish the use of the Hodgkin-Huxley versus the FitzHugh-Nagumo models for the sake of clarity, but we present the detailed sample data in Table 2.

Results: total estimation cost and number of samples. Results are shown in Figure 5, where we show the total cost of each method, and in Table 2, where we report the number of samples taken for each model. For Problem 2 MLBLUE is roughly five times faster than MLMC and two times faster than MFMC. Again, MLBLUE employs a different model selection strategy than the other methods. In fact, MLBLUE is the only method that uses the FitzHugh-Nagumo models and the ODE models, which are much cheaper. These models yield poor approximations to some of the QoIs, causing the model selection procedure (cf. Subsection 5.1.1) of MLMC and MFMC to discard them. In contrast, multi-output MLBLUE includes these models in the estimation and exploits their correlations, thus considerably reducing the total number of high-fidelity samples.

	H-H PDE			F-N PDE	H-H ODE			F-N ODE		
	fine	med.	coarse	coarse	fine	med.	coarse	fine	med.	coarse
MLMC	11959	169150	3856682	0	0	0	0	0	0	0
MFMC	3561	43181	1792145	0	0	0	0	0	0	0
MLBLUE	626	24527	1249805	95652	938	6441	259634	69746	69121	6356360

Table 2: Number of samples taken from each model by the MOSAP-optimal MLBLUE, MLMC and MFMC methods for Problem 2. “H-H” is short for “Hodgkin-Huxley” and “F-N” for “FitzHugh-Nagumo”, while “coarse”, “med.”, and “fine” correspond to the grid and timestep refinement levels $i = 1, 2, 3$ respectively. MLMC and MFMC only use the Hodgkin-Huxley PDE models, while the fine- and medium-grid FitzHugh-Nagumo PDE models are discarded by all methods (hence not shown).

Remark 5.1. While not required for the examples considered in this section, the model selection procedure must account for models that do not admit some of the QoIs as output. For multi-output MLBLUE, this situation is handled automatically by constructing the Ψ_s matrices as described at the beginning of Section 4. Nevertheless, manual model pruning could still be beneficial to restrict the number of possible model groupings, cf. Remark 2.6. For MLMC and MFMC, a simple (though suboptimal) option is to simply discard such models from the estimator. Another option would be to combine multiple single-output MLMC/MFMC estimators within a single joint optimization procedure. We leave the investigation of such a strategy to future work.

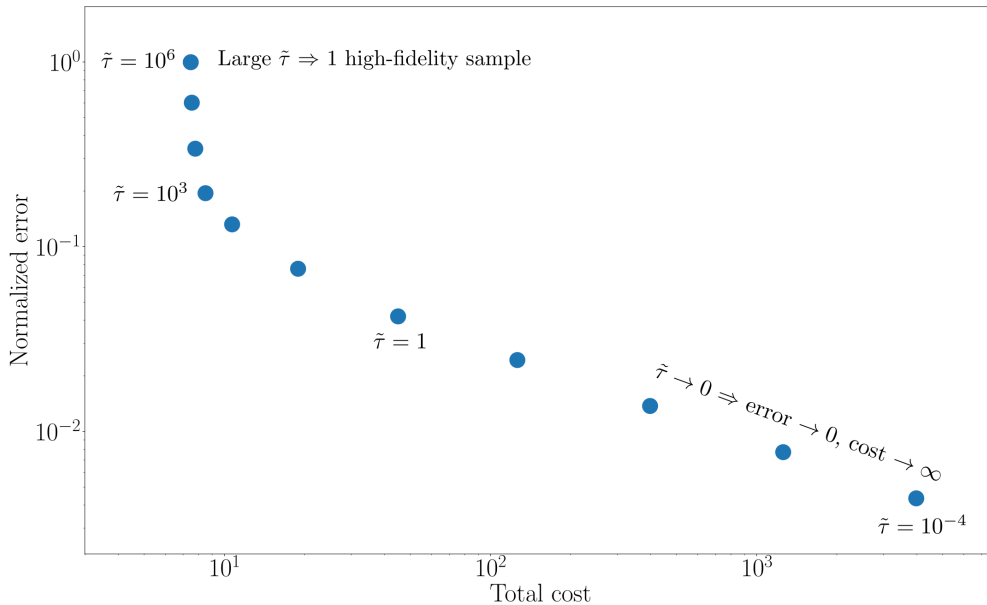


Figure 6: Pareto-front optimization of MLBLUE for Problem 2. Here $\tilde{\tau} = \tau \|\mathbf{c}\|_2$, where τ is the Pareto parameter from (25) and \mathbf{c} is the model cost vector. From right to left, the points in the figure correspond to $\tilde{\tau} = 10^i$ for $i = -4, \dots, 6$.

Results: Pareto frontier. We employ the MOSAP (25) to trace the Pareto frontier of MLBLUE for Problem 2. Results are shown in Figure 6 where we plot the Pareto points corresponding to $\tau = \tilde{\tau} \|\mathbf{c}\|_2^{-1}$, where $\tilde{\tau} = 10^i$ for $i = -4, \dots, 6$. Other than the chosen values of $\tilde{\tau}$, the quantities in the figure are computed in the same way as for Figure 4. The behavior is similar to what we observed for Problem 1 with the exception that the transition to the asymptotic regime appears to be smoother for this problem.

5.1.3. Numerical optimization of the MOSAPs

The objective of this subsection is to benchmark the efficiency and robustness of the new SDP MOSAP formulations against the standard formulations. We consider Problems 1 and 2, and all possible model groupings of size at most χ with $\chi \in \{3, 5, 7\}$ (cf. Remark 2.6).

We solve the SDPs (13), (17), (21) and (23) using the `conelp` solver in the CVXOPT open-source conic optimization software package⁷ [19]. For the standard MOSAP formulations (10), (16), (20), and (22), we use the `trust-constr` trust-region constrained optimization solver from SciPy⁸ and the open-source black-box interior-point optimization solver IPOPT⁹ [26]. We set an absolute error tolerance of 10^{-6} for all optimization solvers. For the budget-constrained MOSAPs we set the budget $b = 10^4 \max(\mathbf{c})$, where \mathbf{c} is the cost vector for all problems. For the statistical-error-constrained MOSAPs we select $\varepsilon_s = 10^{-3} \sqrt{\widehat{\text{V}}[p_1^s]}$, for $s = 1, \dots, m$, where m is the number of QoIs in the problem and $\sqrt{\widehat{\text{V}}[p_1^s]}$ is the estimated standard deviation of the s -th QoI of the high-fidelity model.

We report that the `trust-constr` method fails to converge to the optimal solution in 5000 iterations for both problems and all values of χ . IPOPT does converge, but often fails unless its solver parameters are set accordingly. In particular, we had to set `bound_relax_factor` to 10^{-30} and `honor_original_bounds` to `yes`. The CVXOPT SDP solver always converges in less than 100 iterations. The only problem for which CVXOPT requires a change in the default solver parameters is the budget-constrained MOSAP formulation of Problem 1 for which setting the feasibility tolerance `feastol` to 10^{-3} is needed to ensure convergence.

We show the IPOPT and CVXOPT SDP solver timings in Table 3. The new SDP formulations can robustly be solved in a matter of seconds or minutes depending on the problem size and are much faster to solve than the standard formulations.

Problem	MOSAP	Minimize statistical error			Minimize total cost		
		$\chi = 3$	$\chi = 5$	$\chi = 7$	$\chi = 3$	$\chi = 5$	$\chi = 7$
Problem 1 (12 models, 6 QoIs)	SDP	1.67	66	540	1.2	37	340
	std	1.96	215	2500	3.04	627	5000
Problem 2 (12 models, 4 QoIs)	SDP	1.42	56	518	1.14	44	292
	std	1.54	163	2580	1.85	314	8983

Table 3: CPU time (in seconds) needed to solve the MLBLUE MOSAPs to optimality on a Dell XPS 13 laptop (processor: Intel i7-1185G7 3GHz, RAM: 32GB). The SDP MOSAPs are solved with CVXOPT, while the standard MOSAP formulations are solved with IPOPT.

Remark 5.2. MLBLUE is more sensitive than MLMC and MFMC to integer projections. In fact, we note that taking the ceiling of the real-valued MOSAP solution \mathbf{n}^* is generally a bad idea since \mathbf{n}^* has often many entries that are not close to an integer. A better approach, which we employ in BLUEST, is to take the floor or ceiling of all such entries and try out all feasible combinations. This procedure can be accelerated and made reliable via heuristics that we do not describe here for the sake of brevity. We leave the investigation of a more sophisticated approach, using perhaps a branch and bound algorithm, to future research.

5.2. Coupling restrictions

In this section we consider the third and last test problem. The objective here is to demonstrate that MLBLUE can automatically handle coupling restrictions. In particular, Problem 3 represents a common scenario in scientific computing in which the high-fidelity models are prohibitively expensive and cannot be simulated more than a handful of times. We will also investigate two different approaches for model covariance estimation in this context.

⁷Available at <https://cvxopt.org/>.

⁸Available at <https://scipy.org/>.

⁹Available at <https://coin-or.github.io/Ipopt/>.

5.2.1. *Problem 3: linear random diffusion with restricted high-fidelity samples*

Problem description. As our final test problem we choose a standard linear diffusion partial differential equation (PDE) with random field diffusion coefficient in 2D:

$$\begin{cases} -\nabla \cdot (a(\mathbf{x}, \omega) \nabla u(\mathbf{x}, \omega)) = f(\mathbf{x}), & \mathbf{x} \in \mathcal{D}_1 = (0, 1)^2, \\ u(\mathbf{x}, \omega) = g(\mathbf{x}, \omega), & \mathbf{x} \in \partial \mathcal{D}_1. \end{cases} \quad (38)$$

Here, $f(\mathbf{x}) = \exp(-\|\mathbf{x} - 0.5\|_2^2)$ and $g(\mathbf{x}, \omega) = \theta_1(\omega)x_1^2 + \theta_2(\omega)x_2^2 + 2\theta_3(\omega)x_1x_2$, where $\theta_1, \theta_2, \theta_3$ are independent zero-mean Gaussian random variable with variance 2^{-4} . We set the diffusion coefficient to be $a = e^z$, where $z = z(\mathbf{x}, \omega)$ is a Gaussian field with a Matérn covariance matrix with smoothness parameter $\nu = 1$ and correlation length $\lambda = 0.2$. We choose the mean and standard deviation of z so that $\mathbb{E}[a] = 1.0$, and $\mathbb{V}[a]^{1/2} = 0.3$ for all $\mathbf{x} \in \mathcal{D}_1$. For the definition and properties of Matérn Gaussian fields, we refer the reader to [27]. For sampling the Matérn field, we use the stochastic PDE approach by Lindgren et al. [28] combined with the linear cost-complexity sampling technique by Croci et al. [29, 30].

Quantity of interest. We consider a single QoI given by $p = \int_{\mathcal{D}_1} u^2 dx$.

Model set. We define our model set by simple grid refinement using a hierarchy of $\ell = 7$ uniform triangular grids so that model $i \in \{1, \dots, 7\}$ corresponds to a grid with 2^{8-i} cells in each spatial direction.

Solver and model costs. We solve (38) with the finite element (FE) method using the open-source FE library FEniCS [23] and piecewise-linear Lagrange elements using the conjugate gradient method preconditioned with algebraic multigrid. Cost estimation via CPU timings yields a roughly linear solver cost complexity and we thus define the model sampling costs to be the number of degrees of freedom in each FE model for simplicity.

Coupling restriction. We restrict the number of samples of the two highest-fidelity models to a fixed number $n_{\text{HF}} \in \{2, 4, 8, 16\}$ and investigate how the results vary with n_{HF} . For a study of how MLBLUE performs for a similar diffusion problem, but without any sampling restrictions, we refer the reader to the original papers by Schaden and Ullmann [8, 9]. We remark that their work also includes efficiency comparisons between MLBLUE and other estimators, including MLMC, MFMC and ACV methods, which MLBLUE outperforms.

Covariance estimation. We compare two approaches for estimating covariances.

- a) We take n_{HF} pilot samples of all models and compute sample covariances with these samples.
- b) We take 16 pilot samples of the low-fidelity models, compute their sample covariances, then use extrapolation to approximate the covariances of the high-fidelity models. Specifically:
 1. With the low-fidelity pilot samples, compute all low-fidelity model covariances, as well as the variance of the difference between model pairs, i.e., $\mathbb{V}[p_i - p_j]$ for all $i, j \in \{3, \dots, 7\}$, $i \neq j$.
 2. Draw a single deterministic sample of all models by fixing $a = 1$, $\theta_1 = \theta_2 = \theta_3 = 0$ (their average values), and estimate the convergence rate r with respect to the model degrees of freedom.
 3. Use Richardson extrapolation with rate $2r$ to extrapolate the values of $\mathbb{V}[p_i]$ and $\mathbb{V}[p_i - p_{i+j}]$ for $i = 1, 2$ and $j = 1, \dots, \bar{d}$ from their low-fidelity values. Here \bar{d} is an integer parameter satisfying $1 \leq \bar{d} \leq \ell - 2$ that determines how many covariances are being extrapolated: we estimate the covariances between each high-fidelity model and the next \bar{d} lower-fidelity models. We will be varying \bar{d} in our experiments.
 4. Approximate the high-fidelity model covariances using the formula $C(p_i, p_{i+j}) = \frac{1}{2}(\mathbb{V}[p_i] + \mathbb{V}[p_{i+j}] - \mathbb{V}[p_i - p_{i+j}])$ for $i = 1, 2$ and $j = 1, \dots, \bar{d}$.

In step 2 we use the fact that for PDEs with random coefficients the variance convergence rate is typically twice the deterministic convergence rate, see e.g., [31]. Note that in steps 3-4 we do not necessarily approximate the covariances between all models, hence some model groupings will not be allowed. Nevertheless, restricted model couplings are not problematic for MLBLUEs.

Estimator setup. We proceed as follows. First, we fix a value of $n_{\text{HF}} \in \{2, 4, 8, 16\}$, and either Approach a) or b). Then, we solve the budget-constrained MOSAP (13) with budget given by the cost of sampling all models $2^7 = 128$ times, with maximum allowed model group size $\chi = 4$ (cf. Remark 2.6), and with the additional restriction that the high-fidelity models cannot be sampled more than n_{HF} times.

Results: estimator efficiency. After denoting with θ the set of pilot samples taken, we compute the log of the resulting MLBLUE normalized estimator efficiency:

$$\log_{10} \left(\frac{E(\theta)}{E(\theta_{\text{ex}})} \right) = \log_{10} \left(\frac{\varepsilon^2(\theta_{\text{ex}})}{\mathbb{V}[\hat{\mu}_1]} \right), \quad (39)$$

where $E(\theta)$ is the estimator efficiency defined as

$$E(\theta) = \frac{1}{b \mathbb{V}[\hat{\mu}_1]}, \quad (40)$$

and $\hat{\mu}_1 = \hat{\mu}_1(\theta)$ is the MLBLUE constructed from a single realization θ of the pilot samples, with $\mathbb{V}[\hat{\mu}_1](\theta)$ being its variance. Here $E(\theta_{\text{ex}}) = (b\varepsilon^2(\theta_{\text{ex}}))^{-1}$ and $\varepsilon^2(\theta_{\text{ex}})$ is the variance of a MLBLUE estimator with $n_{\text{HF}} = \infty$ constructed from an exact model covariance estimation using 1000 pilot samples. $E(\theta_{\text{ex}})$ essentially represents the efficiency of the best MLBLUE that can be constructed for this problem with the prescribed budget. We call the quantity in (40) “estimator efficiency” since the inverse product between a Monte Carlo estimator cost and its variance indicates whether an increased budget leads to a more accurate estimator. In fact, the accuracy of any Monte Carlo estimator for an unbiased QoI is proportional to the inverse square root of the cost, i.e., $\varepsilon = O(\text{cost}^{-1/2})$. Therefore, the efficiency $E = (\text{cost } \varepsilon^2)^{-1}$ is always an order-1 quantity in the estimator cost. Nevertheless, a more efficient method will lead to a higher accuracy, a smaller value of ε , and a higher value of E . Note that the closer (39) is to 0, the closer the efficiency of the constructed MLBLUE will be to the best estimator efficiency.

Since poor values of θ may lead to underestimating or overestimating the estimator variance, we draw 50 independent realizations of the estimator to approximate $\mathbb{V}[\hat{\mu}_1](\theta)$. In Figure 7, we study the effect of the high-fidelity sample restrictions and of the model covariance estimation approaches described in Section 5.2.1 across 50 independent realizations of the pilot samples θ . We see how Approach a) leads to poor estimator performance for $n_{\text{HF}} \in \{2, 4\}$, suggesting that this is not a good approach when the high-fidelity samples are severely restricted. Approach b) with $\bar{d} = 1$ performs well in general, but it is not robust to an inaccurate pilot sample estimation. In fact, some realizations of θ lead to a poor estimator efficiency. On the other hand, Approach b) for $\bar{d} > 1$ yields good results for all values of n_{HF} indicating that extrapolation of multiple model cross-covariances, when possible, should be preferred. We report that for both approaches the estimator efficiency does not improve further as n_{HF} increases beyond 16. This result indicates that MLBLUE can handle coupling restrictions with little effect on its efficiency.

Remark 5.3. The values of (39) in Figure 7 that are larger than 0 are likely due to noise in the estimation of $\mathbb{V}[\hat{\mu}_1]$. To produce this experiment we sampled $50^2 = 2500$ independent estimators. Making the estimation of the estimator variance more accurate, while prohibitively expensive, would hardly change the qualitative behavior of these results.

Remark 5.4. A consequence of not knowing all correlations in Approach b) with $\bar{d} = 1$ is that the only allowed model groups that include the two highest-fidelity models \mathcal{M}_1 and \mathcal{M}_2 are

$$\{\mathcal{M}_1\}, \{\mathcal{M}_2\}, \{\mathcal{M}_1, \mathcal{M}_2\}, \{\mathcal{M}_2, \mathcal{M}_3\},$$

i.e., the two models can only be sampled on their own or coupled in pairs. On the other hand, all the groupings between the low-fidelity models $\mathcal{M}_3, \dots, \mathcal{M}_7$ are allowed. This coupling structure could be seen as a way of combining a MLMC estimator on the high-fidelity models (coupled in pairs) with a MFMC estimator on the low-fidelity models (all coupled together), cf. Example 2.1. The MLBLUE algorithm can automatically find such an estimator with little input from the user.

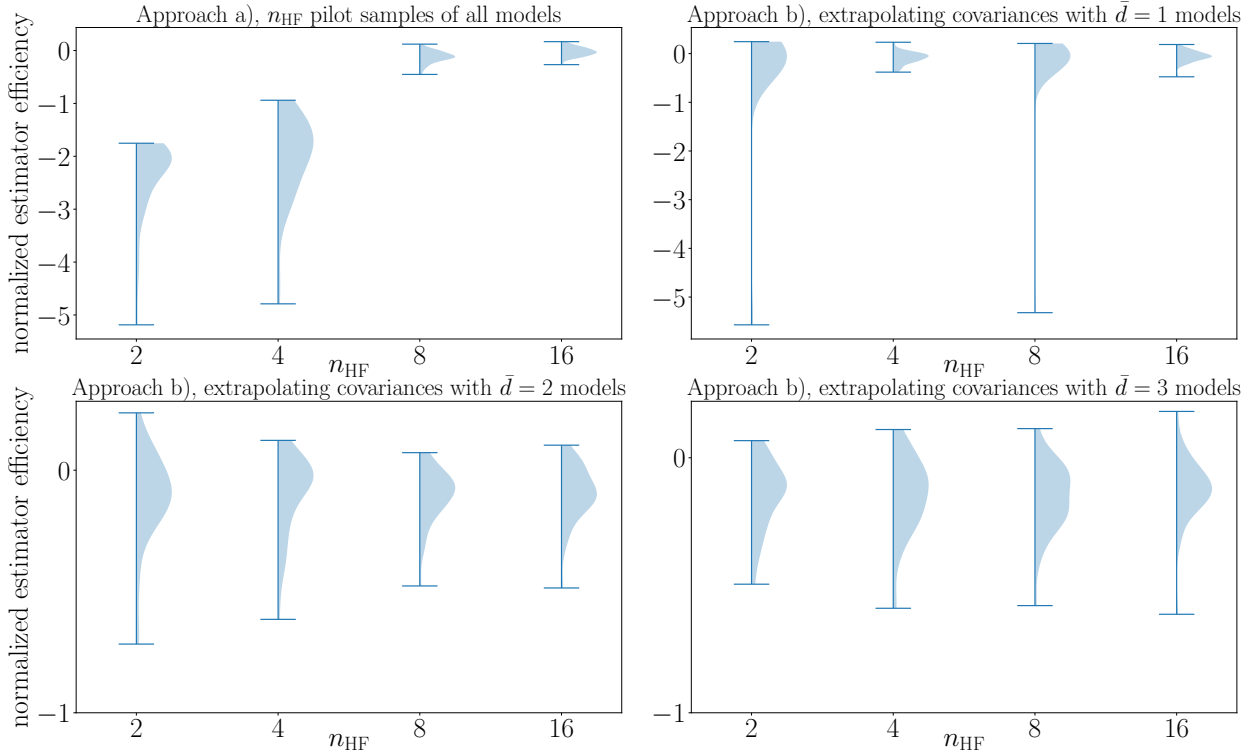


Figure 7: Half-violin plots of the normalized estimator efficiency (39) for the approaches for Problem 3 described in Section 5.2.1 versus the number of restricted high-fidelity samples n_{HF} . The light-blue areas depict kernel-density estimates of the probability density function of (39) with respect to θ . Values close to 0 indicate near-best efficiency.

6. Conclusions

MLBLUEs are optimal across all multilevel and multifidelity linear unbiased estimators and they should thus be preferred in principle over any other such method. However, MLBLUEs (as well as all multilevel/-multifidelity Monte Carlo methods) may lose efficiency if their MOSAP is not solved exactly or if the model covariance estimation is prohibitive.

In this paper we addressed both issues. We formulated the MLBLUE MOSAPs as semi-definite programs that can be solved reliably and efficiently and we showed that these estimators are robust with respect to limited and/or restricted pilot samples. Further, we have extended MLBLUE to the multi-output case by concurrently setting up multiple estimators for all QoIs, thus making MLBLUEs the only methods with an optimal strategy for handling multiple outputs. Consequently, this work enables the fast, reliable design and setup of MLBLUEs for forward UQ problems involving one or more QoIs. Numerical experimentation shows the resulting MLBLUEs to be more efficient than MLMC and MFMC. MLBLUEs can also deal efficiently with heterogeneous model sets.

Additional work is still needed to design a comprehensive MLBLUE algorithmic framework: Integer programming strategies for the MOSAP could be beneficial since the continuous relaxation of the MOSAP may be far from optimal for some problems (cf. Remark 5.2). Furthermore, specific strategies should be designed to decrease the cost of offline estimation and/or allow pilot sample reuse. For instance, QoIs with large kurtosis require a large number of pilot samples [4] which, if re-used within MLBLUE, would make the estimator biased [8], and hence not necessarily optimal. Finally, MLBLUE advanced estimation strategies for rare-event simulation, nested expectations, and non-smooth QoIs are yet to be investigated.

While the most efficient method for a given application is, in general, problem-dependent, single- and multi-output MLBLUEs are optimal whenever the above challenges are not relevant and should thus be the

methods of choice. MLBLUE is currently the only method with an optimal multi-output strategy and that is capable of automatically addressing model selection, making it a powerful addition to the forward UQ toolbox.

Acknowledgements

MC thanks Nicole Aretz for her helpful feedback on Section 5. MC and KW's research is based upon work supported by the Department of Energy NNSA Award Number DE-NA0003969. SW acknowledges support from a J. T. Oden Faculty Fellowship from the Oden Institute at the University of Texas at Austin as well as NSF Awards DMS 2023239 and CCF 2224213, DOE via subcontract 8F-30039 from Argonne National Laboratory. KW and SW additionally acknowledge support from the AFOSR MURI on Machine Learning of Physics-based Systems, AFOSR Grant FA9550-21-1-0084.

Declaration of competing interest

The authors declare that they have no competing interests.

References

- [1] S. Heinrich, Multilevel Monte Carlo methods, in: *International Conference on Large-Scale Scientific Computing*, Springer, 2001, pp. 58–67.
- [2] M. B. Giles, Multilevel Monte Carlo path simulation, *Operations Research* 56 (3) (2008) 607–617.
- [3] L. W. T. Ng, K. E. Willcox, Multifidelity approaches for optimization under uncertainty, *International Journal for Numerical Methods in Engineering* 100 (10) (2014) 746–772.
- [4] M. B. Giles, Multilevel Monte Carlo methods, *Acta Numerica* 24 (2015) 259–328.
- [5] A.-L. Haji-Ali, F. Nobile, R. Tempone, Multi-index Monte Carlo: when sparsity meets sampling, *Numerische Mathematik* 132 (4) (2016) 767–806.
- [6] B. Peherstorfer, K. Willcox, M. Gunzburger, Optimal model management for multifidelity Monte Carlo estimation, *SIAM Journal on Scientific Computing* 38 (5) (2016) A3163–A3194.
- [7] A. A. Gorodetsky, G. Geraci, M. S. Eldred, J. D. Jakeman, A generalized approximate control variate framework for multifidelity uncertainty quantification, *Journal of Computational Physics* 408 (2020) 109257.
- [8] D. Schaden, E. Ullmann, On multilevel best linear unbiased estimators, *SIAM/ASA Journal on Uncertainty Quantification* 8 (2) (2020) 601–635.
- [9] D. Schaden, E. Ullmann, Asymptotic analysis of multilevel best linear unbiased estimators, *SIAM/ASA Journal on Uncertainty Quantification* 9 (3) (2021) 953–978.
- [10] L. Vandenberghe, S. Boyd, Semidefinite programming, *SIAM Review* 38 (1) (1996) 49–95.
- [11] S. J. Benson, Y. Ye, DSDP5: Software for semidefinite programming, Tech. Rep. ANL/MCS-P1289-0905, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL (2005).
- [12] R. H. Tütüncü, K.-C. Toh, M. J. Todd, Solving semidefinite-quadratic-linear programs using SDPT3, *Mathematical Programming* 95 (2) (2003) 189–217.
- [13] E. D. Andersen, K. D. Andersen, The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm, in: *High Performance Optimization*, Springer, 2000, pp. 197–232.
- [14] S. Diamond, S. Boyd, CVXPY: A Python-embedded modeling language for convex optimization, *The Journal of Machine Learning Research* 17 (1) (2016) 2909–2913.
- [15] B. Peherstorfer, K. Willcox, M. Gunzburger, Survey of multifidelity methods in uncertainty propagation, inference, and optimization, *SIAM Review* 60 (3) (2018) 550–591.
- [16] J. Groß, The general Gauss-Markov model with possibly singular dispersion matrix, *Statistical Papers* 45 (3) (2004) 311–336.
- [17] C. R. Rao, H. Toutenburg, H. C. Shalabh, M. Schomaker, *Linear models and generalizations*, in: *Least Squares and Alternatives* (3rd edition), Springer, Berlin Heidelberg New York, 2008.
- [18] S. Boyd, L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2003.
- [19] M. Andersen, J. Dahl, Z. Liu, L. Vandenberghe, *Interior-point methods for large-scale cone programming*, MIT Press, Cambridge, Mass., 2012, Ch. 3, pp. 55–83.
- [20] G. Zyskind, F. B. Martin, On best linear estimation and general Gauss-Markov theorem in linear models with arbitrary nonnegative covariance structure, *SIAM Journal on Applied Mathematics* 17 (6) (1969) 1190–1202.
- [21] A. Albert, Conditions for positive and nonnegative definiteness in terms of pseudoinverses, *SIAM Journal on Applied Mathematics* 17 (2) (1969) 434–440.
- [22] F. Alizadeh, D. Goldfarb, Second-order cone programming, *Mathematical Programming* 95 (1) (2003) 3–51.
- [23] A. Logg, K.-A. Mardal, G. Wells, *Automated Solution of Differential Equations by the Finite Element Method: the FEniCS Book*, Vol. 84, Springer Science & Business Media, 2012.

- [24] S. J. Chapman, A. C. Fowler, R. Hinch, S. L. Waters, *An Introduction to Mathematical Physiology*, Mathematical Institute, The University of Oxford, 2021.
- [25] H. Lieberstein, On the Hodgkin-Huxley partial differential equation, *Mathematical Biosciences* 1 (1) (1967) 45–69.
- [26] A. Wächter, L. T. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, *Mathematical Programming* 106 (1) (2006) 25–57.
- [27] P. Abrahamsen, *A Review of Gaussian Random Fields and Correlation Functions*, 2nd Edition, Norwegian Computing Center, 1997.
- [28] F. Lindgren, H. Rue, J. Lindström, An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73 (4) (2009) 423–498.
- [29] M. Croci, M. B. Giles, M. E. Rognes, P. E. Farrell, Efficient white noise sampling and coupling for multilevel Monte Carlo with non-nested meshes, *SIAM/ASA Journal on Uncertainty Quantification* 6 (4) (2018) 1630–1655. [doi:10.1137/18M1175239](https://doi.org/10.1137/18M1175239).
- [30] M. Croci, M. Giles, P. E. Farrell, Multilevel quasi Monte Carlo methods for elliptic PDEs with random field coefficients via fast white noise sampling, *SIAM Journal on Scientific Computing* 43 (4) (2021) A2840–A2868.
- [31] K. A. Cliffe, M. B. Giles, R. Scheichl, A. L. Teckentrup, Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients, *Computing and Visualization in Science* 14 (1) (2011) 3–15. [doi:10.1007/s00791-011-0160-x](https://doi.org/10.1007/s00791-011-0160-x).