Journal of Computational Physics 388 (2019) 300–315

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

High-dimensional and higher-order multifidelity Monte Carlo estimators

A. Quaglino *, S. Pezzuto *, R. Krause

Center for Computational Medicine in Cardiology, Institute of Computational Science, Università della Svizzera italiana, Lugano, Switzerland

ARTICLE INFO

Article history:
Received 27 July 2018
Received in revised form 25 February 2019
Accepted 19 March 2019
Available online 21 March 2019

Keywords:
Uncertainty quantification
Monte Carlo
Multifidelity
Global sensitivity analysis
Model reduction
Cardiac electrophysiology

ABSTRACT

Multifidelity Monte Carlo methods rely on a hierarchy of possibly less accurate but statistically correlated simplified or reduced models, in order to accelerate the estimation of statistics of high-fidelity models without compromising the accuracy of the estimates. This approach has recently gained widespread attention in uncertainty quantification [1]. This is partly due to the availability of optimal strategies for the estimation of the expectation of scalar quantities-of-interest [2]. In practice, the optimal strategy for the expectation is also used for the estimation of variance and sensitivity indices [3]. However, a general strategy is still lacking for vector-valued problems, nonlinearly statistically-dependent models, and estimators for which a closed-form expression of the error is unavailable. The focus of the present work is to generalize the standard multifidelity estimators to the above cases. The proposed generalized estimators lead to an optimization problem that can be solved analytically and whose coefficients can be estimated numerically with few runs of the high- and low-fidelity models. We analyze the performance of the proposed approach on a selected number of experiments, with a particular focus on cardiac electrophysiology, where a hierarchy of physics-based low-fidelity models is readily available.

© 2019 Elsevier Inc. All rights reserved.

1. Introduction

Sampling methods for Uncertainty Quantification (UQ) require multiple evaluations of the problem at hand. Often, each sample requires computing the solution of partial differential equations (PDEs), which is generally a computationally demanding task. For example, in cardiac electrophysiology, a single patient-tailored simulation based on the bidomain model can take thousands of node-hours on a large cluster [4]. Therefore, UQ for such models at organ scale is currently unfeasible with plain Monte Carlo methods, although its importance has been highlighted in the literature [5]. A more sophisticated approach is to perform most of the simulations on a hierarchy of low-resolution models, yielding the so-called (geometric) multilevel Monte Carlo (MLMC) method [6], which extends the idea of control variates methods [7]. The use of the highest-resolution level guarantees convergence, while a significant portion of the computational load is offset to the low-resolution hierarchy. While this can offer a significant speedup of the estimation, it is not always possible to coarsen complex geometries and to robustly transfer information from one level to another, to an extent that the full potential of this approach can be reached [8]. This is for example the case in the context of cardiac modeling [9]. Moreover, the estimation of high-order moments requires special care [10,11].

* Corresponding authors.
E-mail addresses: alessio.quaglino@usi.ch (A. Quaglino), simone.pezzuto@usi.ch (S. Pezzuto).

https://doi.org/10.1016/j.jcp.2019.03.026
0021-9991/© 2019 Elsevier Inc. All rights reserved.

An alternative approach, which does not rely on geometry coarsening, is to use model reduction strategies, such as projection-based or surrogate models [12,13]. Such an approximation is build upon the observation that in many applications it is clearly possible to distinguish between *offline* and *online* phases of the workflow, where the former typically can be very expensive and encompass everything that can be precomputed in advance, e.g., training a surrogate model before any patient-specific data becomes available, while the latter should be very cheap and consist only of evaluations of the reduced model. Unfortunately, this approach has two significant drawbacks. On the one hand, complex error estimates must be provided to ensure that the approximation error is within acceptable bounds. On the other hand, repeating the offline training is very expensive but may often be necessary, e.g., in the case of patient-specific UQ. Alternative to that, physics-based reduction strategies are also possible: in cardiac electrophysiology, the activation map can readily be computed via the eikonal equation, which provides a physiologically-meaningful solution [14], strongly linked to the bidomain equation [15]. Error bounds, however, are not trivially found, limiting the applicability of the methodology.

A more recent idea is to replace the above paradigm of model *selection* with that of model *fusion*, where low-fidelity models are allowed to be inaccurate, in the sense of not providing an approximation within certain error bounds, provided that they exhibit some degree of correlation [2] or even just statistical dependence [16] to the high-fidelity one. This approach has a twofold advantage. Firstly, it is the statistical dependence, rather than the error bounds of coarse models, that is crucial to ensure that propagating uncertainties via the low-fidelity models provides useful information on the statistics of the high-fidelity quantity-of-interest. Secondly, in the case of PDEs, low-fidelity models are not restricted to low-resolution geometries, therefore solving them can be several orders of magnitude faster than a coarse model. If a statistical dependence between the models is present, in practice, most of the computational effort is taken by the estimation of the linear correlation coefficients or a nonlinear map between the outputs of the low- and the high-fidelity models. For such reasons, multifidelity Monte Carlo (MFMC) methods have become very popular over the last years and their applications span the fields of UQ, inverse problems, and optimization [1].

Several approaches can be used to combine the outputs of the model hierarchy. The use of a Bayesian regression, mapping the low-fidelity output to the high-fidelity one, has been advocated in [16]. In [17], a Gaussian Process Regression is considered in the context of cardiac electrophysiology. The Bayesian nature of this approach automatically augments the estimate with full probability distributions and credible intervals, obtained independently of the degree of statistical dependence between models and the number of samples employed. This makes it a very suitable approach for scenarios where resources are scarce, such as clinical applications. In contrast to that, the multifidelity Monte Carlo method [2] provides point estimates, rather than probability distributions. However, these are equipped with convergence estimates for the asymptotic case, i.e., their error vanishes as the number of samples goes to infinity. This makes this approach particularly interesting when computational resources are large (but finite) and accurate estimates are desired.

Current MFMC methods have been successfully used for estimating the expectation, the variance, the Sobol indices, and rare events [1]. For the expectation of a scalar quantity-of-interest, the optimal multifidelity estimator for a given computational budget can be found analytically [2]. Similarly, Qian et al. [3] proposed an estimator for variance and Sobol indices, with an analytical derivation only for the former. Interestingly, the authors also report that the optimal weights and sampling strategy for the expectation multifidelity estimator performs very well also for the variance and Sobol indices, although not being optimal. This observation suggests that a multifidelity estimator for a general statistic of interest may perform relatively well even when the weights and sampling strategy are not optimal.

In this work, we explore the above idea, hence extending the aforementioned multifidelity estimators in several ways. First, we generalize the scalar-valued multifidelity estimator to vector-valued quantity-of-interests, proposing an optimal estimator in terms of mean squared error (MSE) that closely resembles the one obtained in the scalar case. Second, we analyze a broad class of high-order statistics of interest, for which a closed-form expression of the error is unavailable and therefore the optimal sampling strategy is unknown. To overcome this limitation, we suggest to approximate the MSE by estimating the correlations between the estimators, via sampling of the low- and high-fidelity models. Third, we suggest a strategy to leverage nonlinear statistical dependence between the models, which is closely connected to [16,17].

The paper is organized as follows: Section 2 introduces the standard multifidelity Monte Carlo estimators of scalar quantities-of-interest. Section 3 describes the proposed generalized multifidelity estimators. Section 4 presents the results of the numerical experiments for the assessment of each estimator. Section 5 details a real-world application to cardiac electrophysiology. Section 6 illustrates an application to nonlinear incompressible elasticity.

2. Standard multifidelity Monte Carlo estimators

Let upper-case letters denote random variables (e.g. $\mathbf{S}$, Ψ) and lower-case letters the values these take (e.g. $\mathbf{s}$, ψ), where ψ and $\psi(\mathbf{s})$ denote both, the value of the quantity-of-interest (QoI) as well as the function that provides the QoI with respect to the random input $\mathbf{s}$. Let $\psi^{(i)}(\mathbf{s})$ be the i-th model of the multifidelity hierarchy, with $i = 1, \dots, K$. For ease of notations, we denote by $\psi_h(\mathbf{s}) := \psi^{(1)}(\mathbf{s})$ and $\psi_l(\mathbf{s}) := \psi^{(K)}(\mathbf{s})$ the high- and lowest-fidelity models.

The main task in the context of uncertainty propagation is the computation of integrals. In particular, the sought expectation of the QoI is obtained by:

$$\mathrm{E}[\Psi_h] = \int_{\Omega} \psi_h(\mathbf{s})\ p(\mathbf{s})\ \mathrm{d}\mathbf{s}, \tag{1}$$

where $p(\mathbf{s})$ denotes the probability density of the random inputs. Similarly, one can compute other statistical indicators (e.g. moments, densities etc.), by evaluating integrals of the form $\int h(\psi_h(\mathbf{s}))\ p(\mathbf{s})ds$. The method of choice in high-dimensional settings is direct Monte Carlo, where the expectation of Equation (1) is estimated by:

$$\hat{\Psi}_m^{(i)} = \frac{1}{m}\sum_{j=1}^{m} \psi^{(i)}(\mathbf{s}_j), \tag{2}$$

where $i = 1$ and the $\mathbf{s}_j$ are independent, identically distributed samples drawn from $p(\mathbf{s})$. The convergence rate is independent of the dimension of $\mathbf{S}$ and the error decays as $\mathcal{O}(\frac{1}{\sqrt{m}})$ [18]. Therefore, in problems where each evaluation of $\psi_h(\mathbf{s})$ poses a significant computational burden, the use of direct Monte Carlo can become impractical or even infeasible.

2.1. Scalar multifidelity expectation estimator

At this stage, a multifidelity estimator is introduced

$$\hat{\Psi}_h = \hat{\Psi}_{m_1}^{(1)} + \sum_{i=2}^{K} \alpha_i \left(\hat{\Psi}_{m_i}^{(i)} - \hat{\Psi}_{m_{i-1}}^{(i)} \right). \tag{3}$$

The key aspect of the multifidelity estimate is that the magnitude of the discrepancy $|\psi_h(\mathbf{s}) - \psi^{(i)}(\mathbf{s})|$ does not affect its accuracy. More precisely, let the Pearson correlation coefficient of the i-th model be

$$\rho_{i,j} := \frac{\mathrm{Cov}[\psi^{(i)}, \psi^{(j)}]}{\sigma_i \sigma_j} \tag{4}$$

and its variance $\sigma_i^2 := \mathrm{Var}[\hat{\Psi}^{(i)}]$. Then, it is possible to show [2] that the MSE of the multifidelity estimator is

$$\mathrm{Var}[\hat{\Psi}_h] = \frac{\sigma_1^2}{m_1} + \sum_{i=2}^{K} \left(\frac{1}{m_{i-1}} - \frac{1}{m_i} \right) \left(\alpha_i^2 \sigma_i^2 - 2\alpha_i \rho_{1,i} \sigma_1 \sigma_i \right). \tag{5}$$

Moreover, the cost of the estimator is

$$C(\hat{\Psi}_h) = \sum_{i=1}^{K} w_i m_i, \tag{6}$$

where w_i is the computational cost of the i-th model. Under some mild assumptions on $\rho_{i,j}$ and w_i, it is possible to perform a constrained minimization of the variance (5) with respect to α and m, for a prescribed computational cost $C(\hat{\Psi}_h) = B > 0$, obtaining

$$\alpha_i^* := \frac{\rho_{1,i}\sigma_1}{\sigma_i}, \quad m_1^* = \frac{B}{\sum_{i=1}^{K} w_i r_i}, \quad m_i^* = m_1^* r_i, \quad i = 2, \dots, K \tag{7}$$

where

$$r_i := \sqrt{\frac{w_1 \left(\rho_{1,i}^2 - \rho_{1,i+1}^2 \right)}{w_i \left(1 - \rho_{1,2}^2 \right)}}. \tag{8}$$

It follows that the ratio of errors (i.e. the variance reduction) of the standard and the multifidelity estimators, for the same computational budget, is

$$\frac{E(\hat{\Psi}_h^*)}{E(\hat{\Psi}_n^{(1)})} = \left(\sum_{i=1}^{K} \sqrt{\frac{w_i}{w_1} \left(\rho_{1,i}^2 - \rho_{1,i+1}^2 \right)} \right)^2, \tag{9}$$

where $n = B/w_1$. Therefore, the variance is reduced if the costs w_i and the differences of the squared correlation coefficients are low. This condition replaces the classical assumption on the deterministic pointwise errors $|\psi_h(\mathbf{s}) - \psi^{(i)}(\mathbf{s})|$ of the low-fidelity models with respect to the high-fidelity one.

2.2. Variance and Sobol indices estimators

The standard approach to develop multifidelity estimators for higher-order moments $q(\Psi_h)$ is to substitute $\hat{\Psi}_m^{(i)}$ in (3) with the unbiased single-level estimator $\hat{q}$. For example, in the case of variance and Sobol indices, the use of the following estimators is proposed in [3]

$$\hat{V} = \frac{1}{m-1}\sum_{i=1}^{m}\left(\psi(s_i) - \hat{\Psi}\right)^2,$$

$$\hat{V}_j = \frac{2}{2m-1}\left(\sum_{i=1}^{m}\psi(s_i)\psi(y_i^j) - m\left(\frac{\hat{\Psi}+\hat{\Psi}'}{2}\right)^2 + \frac{\hat{V}+\hat{V}'}{4}\right),$$

$$\hat{T}_j = \frac{1}{2m}\sum_{i=1}^{m}\left(\psi(s_i') - \psi(y_i^j)\right)^2,$$

where s' is a second set of samples and y_i^j is equal to s_i', except for the j-th component which is equal to that of s_i. However, obtaining error estimates in the form of (5) is generally too difficult or impractical. Therefore, it is advocated in [3] to employ the optimal strategy (m^*, α^*) for the expectation also for the computation of other estimators. In the following section, we propose a different approach based on estimating the correlations between each pair of summands (as computed via two different models) in the above estimators.

3. Generalized multifidelity Monte Carlo estimators

3.1. Vector-valued multifidelity expectation estimator

We now consider the case of a vector-valued quantity-of-interest. In particular, we focus on a spatially-dependent ψ, i.e. $\psi(x)$. It follows that the random variables and their estimators are also spatially dependent, i.e., $\Psi(x)$ and $\hat{\Psi}(x)$. We consider the finite-dimensional case of computing them at a given set of points $\{x_j\}_{j=1}^N$. Given a computational budget B, the optimal scalar-valued estimator can therefore be computed at each point x_j. By prescribing a tolerance ϵ^2 for the integrated MSE of the optimal estimator, expressed as a function of the computational budget B, as given in [2], we obtain

$$\int_\Omega \frac{\sigma_1^2(x)}{B}\left(\sum_{i=1}^{K}\sqrt{\frac{w_i}{w_1}\left(\rho_{1,i}^2(x) - \rho_{1,i+1}^2(x)\right)}\right)^2 dx < \epsilon^2.$$

By assuming a spatially-uniform budget $B(x) = \overline{B}$, linearity of the integral can be exploited, so that the optimal budget $\overline{B}^*$ that satisfies the above equation is easily found. However, this will yield a spatially-dependent number of model evaluations $m_i(x)$. Since it is not possible to perform a model evaluation only for a subset of the points, the following choices are made

$$m_i = \max_{j=1,\dots,N} m_i(x_j).$$

Unfortunately, in general this will lead to a large number of model evaluations. Instead, we define the integrated vector error as

$$\int_\Omega E(\hat{\Psi}_h, x, \alpha)\, dx := \sum_{j=1}^{N} E(\hat{\Psi}_h, x_j, \alpha)|\Omega_j|, \tag{10}$$

where $|\Omega_j|$ is the measure of the local integration area and

$$E(\hat{\Psi}_h, x_j, \alpha) = \frac{\sigma_1^2(x_j)}{m_1} + \sum_{i=2}^{K}\left(\frac{1}{m_{i-1}} - \frac{1}{m_i}\right)\left(\alpha_i^2(x_j)\sigma_i^2(x_j) - 2\alpha_i(x_j)\rho_{1,i}(x_j)\sigma_1(x_j)\sigma_i(x_j)\right).$$

Therefore, the error of the vector-valued case can be seen as a weighted average of the scalar-valued errors. Hence, the resulting estimator can be also used for the generic spatially-independent vector-valued case, given some importance weights $|\Omega_j|$. Setting the partial derivative of $E(\hat{\Psi}_h, x_j, \alpha)$ with respect to $\alpha_i(x_j)$ equal to zero, we obtain that

$$\alpha_i^*(x_j) := \frac{\rho_{1,i}(x_j)\sigma_1(x_j)}{\sigma_i(x_j)}, \quad \forall j = 1\cdots N.$$

It follows that

$$E(\hat{\Psi}_h, x_j, \alpha^*) = \frac{\sigma_1^2(x_j)}{m_1} - \sum_{i=2}^{K} \left(\frac{1}{m_{i-1}} - \frac{1}{m_i} \right) \rho_{1,i}^2(x_j) \sigma_1^2(x_j).$$

Therefore, the error sum can be written as

$$\sum_{j=1}^{N} E(\hat{\Psi}_h, x_j, \alpha) = \frac{1}{m_1} \overline{\sigma}_1^2 - \sum_{i=2}^{K} \left(\frac{1}{m_{i-1}} - \frac{1}{m_i} \right) \overline{\rho}_{1,i}^2 \overline{\sigma}_1^2, \tag{11}$$

where

$$\overline{\sigma}_1^2 := \sum_{j=1}^{N} \sigma_1^2(x_j) |\Omega_j|, \quad \overline{\rho}_{1,i}^2 := \frac{1}{\overline{\sigma}_1^2} \left(\sum_{j=1}^{N} \rho_{1,i}^2(x_j) \sigma_1^2(x_j) |\Omega_j| \right)$$

Equation (11) has the same form as for the scalar case, hence it leads to

$$m_1^* = \frac{B}{\sum_{i=1}^{K} w_i \overline{r}_i}, \quad m_i^* = m_1^* \overline{r}_i, \quad \overline{r}_i := \sqrt{\frac{w_1 \left(\overline{\rho}_{1,i}^2 - \overline{\rho}_{1,i+1}^2 \right)}{w_i \left(1 - \overline{\rho}_{1,2}^2 \right)}}. \tag{12}$$

3.2. Higher-order and sensitivity estimators

In general, given a single-level estimator $\hat{q}$, its corresponding multifidelity estimator is defined as

$$\hat{q}_h = \hat{q}_{m_1}^{(1)} + \sum_{i=2}^{K} \alpha_i \left(\hat{q}_{m_i}^{(i)} - \hat{q}_{m_{i-1}}^{(i)} \right). \tag{13}$$

Regardless of the true expression of $\hat{q}$, we implicitly define a function q via the relation

$$\hat{q}_{m_i}^{(j)} = \frac{1}{m_i} \sum_{n=1}^{m_i} q(\Psi^{(j)}(s_n)). \tag{14}$$

Under the assumption that the s_n are i.i.d. and $m_{i-1} < m_i$, it can be shown that

$$\mathrm{Cov}[\hat{q}_{m_i}^{(j)}, \hat{q}_{m_{i-1}}^{(k)} - \hat{q}_{m_i}^{(k)}] = \frac{1}{m_i} \left(1 - \min \left(1, \frac{m_i}{m_{i-1}} \right) \right) c_{jk} = 0, \quad \forall (j,k),$$

where $c_{jk} = \mathrm{Cov}[q(\Psi^{(j)}), q(\Psi^{(k)})]$. Therefore, it possible to write the MSE of (13) as

$$\begin{aligned} \mathrm{Var}[\hat{q}_h] = \mathrm{Var}[\hat{q}_{m_1}^{(1)}] &+ \sum_{i=2}^{K} \alpha_i^2 \left(\mathrm{Var}[\hat{q}_{m_{i-1}}^{(i)}] - \mathrm{Var}[\hat{q}_{m_i}^{(i)}] \right) \\ &- 2 \sum_{i=2}^{K} \alpha_i \left(\mathrm{Cov}[\hat{q}_{m_1}^{(1)}, \hat{q}_{m_{i-1}}^{(i)}] - \mathrm{Cov}[\hat{q}_{m_1}^{(1)}, \hat{q}_{m_i}^{(i)}] \right). \end{aligned}$$

While in general it is not possible to find analytic expressions for the above terms, they can be estimated numerically in a preprocessing step by the definition of q in equation (14) as

$$\mathrm{Cov}[\hat{q}_{m_i}^{(j)}, \hat{q}_{m_{i-1}}^{(k)}] \approx \frac{1}{m_i} \frac{1}{N-1} \sum_{n=1}^{N} \left(q(\Psi^{(j)}(s_n)) - \hat{q}_N^{(j)} \right) \left(q(\Psi^{(k)}(s_n)) - \hat{q}_N^{(k)} \right), \quad m_{i-1} \leq m_i,$$

where N is a small number (i.e. $N \approx 10 - 20$). By defining

$$\left(\sigma_i^q \right)^2 := \frac{1}{N-1} \sum_{n=1}^{N} \left(q(\Psi^{(i)}(s_n)) - \hat{q}_N^{(i)} \right)^2,$$

$$\rho_{1,i}^q := \frac{1}{\sigma_1^q \sigma_i^q (N-1)} \sum_{n=1}^{N} \left(q(\Psi^{(1)}(s_n)) - \hat{q}_N^{(1)} \right) \left(q(\Psi^{(i)}(s_n)) - \hat{q}_N^{(i)} \right),$$

we obtain an expression similar to (5)

$$\mathrm{Var}[\hat{q}_h] = \frac{(\sigma_1^q)^2}{m_1} + \sum_{i=2}^{K} \left(\frac{1}{m_{i-1}} - \frac{1}{m_i} \right) \left(\alpha_i^2 (\sigma_i^q)^2 - 2\alpha_i \rho_{1,i}^q \sigma_1^q \sigma_i^q \right). \tag{15}$$

Therefore, the optimization procedure presented above leads to a similar result with σ^q and ρ^q in place of σ and ρ.

3.3. Nonlinearly statistically-dependent estimators

We now consider the case where the correlation coefficient $\rho_{1,i}$ between the models is low, but the two model exhibit a high degree of nonlinear statistical dependence. For ease of notation, we consider to estimate the expectation of a scalar Ψ, but the same procedure can be applied to estimate any statistics and to the vector-valued case, by combining it with the strategies from the previous sections. Let us consider an estimator of the form

$$\hat{g}_h = \hat{\Psi}_{m_1}^{(1)} + \sum_{i=2}^{K} \alpha_i \left(\hat{g}_{m_i}^{(i)} - \hat{g}_{m_{i-1}}^{(i)} \right), \tag{16}$$

where

$$\hat{g}_{m_i}^{(j)} = \frac{1}{m_i} \sum_{n=1}^{m_i} g(\Psi^{(j)}(s_n)). \tag{17}$$

The function g is an additional 1-dimensional model, for example a Gaussian Process regression, which is fitted using N_g samples. The idea is that for a strongly nonlinear statistical dependence, $N_g + N_\rho \ll N$, where N and N_ρ are the number of samples used to estimate the correlation for, respectively, the estimators (3) and (16). Following the same procedure as in the previous section, it follows that

$$\mathrm{Var}[\hat{g}_h] = \frac{(\sigma_1)^2}{m_1} + \sum_{i=2}^{K} \left(\frac{1}{m_{i-1}} - \frac{1}{m_i} \right) \left(\alpha_i^2 (\sigma_i^g)^2 - 2\alpha_i \rho_{1,i}^g \sigma_1 \sigma_i^g \right). \tag{18}$$

Therefore, the optimization procedure presented above leads to a similar result with σ^q and ρ^q in place of σ and ρ, where

$$\left(\sigma_i^g\right)^2 := \frac{1}{N-1} \sum_{n=1}^{N} \left(g(\Psi^{(i)}(s_n)) - \hat{g}_N^{(i)} \right)^2,$$

$$\rho_{1,i}^g := \frac{1}{\sigma_1 \sigma_i^g (N-1)} \sum_{n=1}^{N} \left(\Psi^{(1)}(s_n) - \hat{\Psi}_N^{(1)} \right) \left(g(\Psi^{(i)}(s_n)) - \hat{g}_N^{(i)} \right).$$

3.4. Choosing the computational budget

In order to choose the budget, we express the integrated MSE of the optimal estimator as a function of the computational budget B, as given in [2]

$$E(\hat{\Psi}_h^*, B) = \frac{\overline{\sigma}_1^2}{B} \left(\sum_{i=1}^{K} \sqrt{\frac{w_i}{w_1} \left(\overline{\rho}_{1,i}^2 - \overline{\rho}_{1,i+1}^2 \right)} \right)^2 < \epsilon^2,$$

from which it follows that

$$\overline{B}^* = \left(\frac{\overline{\sigma}_1}{\epsilon} \sum_{i=1}^{K} \sqrt{\frac{w_i}{w_1} \left(\overline{\rho}_{1,i}^2 - \overline{\rho}_{1,i+1}^2 \right)} \right)^2. \tag{19}$$

4. Numerical assessment

In this section, we test the proposed generalized estimators on two simple benchmark problems. The first test is a well-known problem in the context of multifidelity methods [3]. There, we consider the expectation, variance, and Sobol index estimators. In the second test, we focus on comparing the linear and nonlinearly statistically-dependent estimators. In Figs. 3 and 6, the two approaches are referred to as Linear and GPR, respectively. All of the multifidelity routines are implemented in a in-house, newly developed Python library called SLOTH.

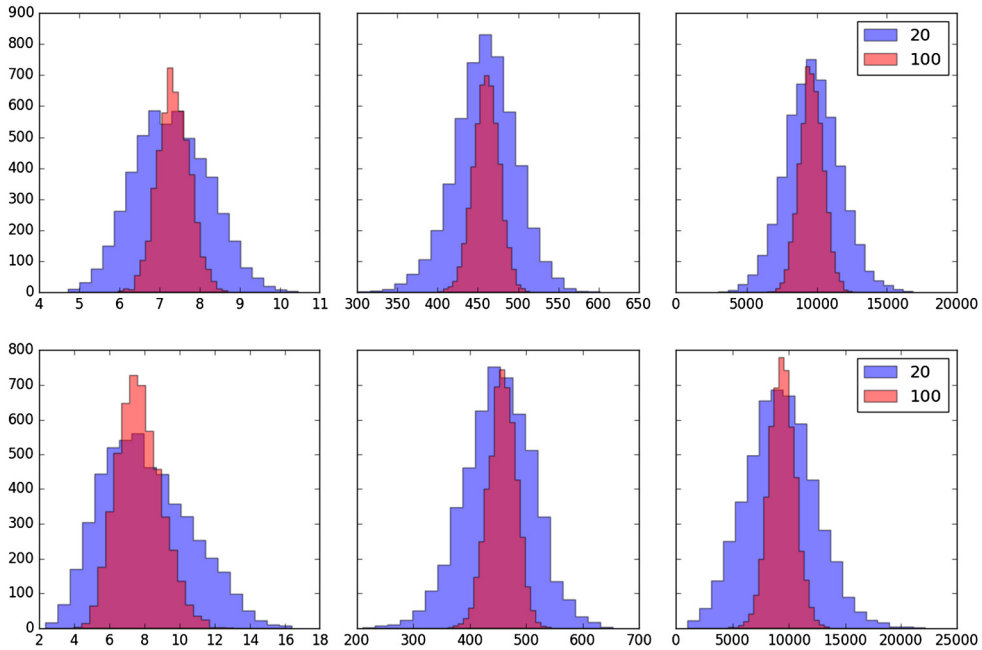


Fig. 1. Distribution of optimal m_k for the expectation (top row) and variance (bottom row) estimators for the Ishigami model with 20 (blue) and 100 (red) preprocessing samples, computed with 5000 replicates. Columns: m_1, m_2, and m_3. Plot axes: values (horizontal) and frequency (vertical) of the estimator. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

4.1. Ishigami model

We consider the Ishigami model [19] and the low-fidelity models proposed in [3]

$$f^{(1)} = \sin(z_1) + 5\sin^2(z_2) + \frac{1}{10} z_3^4 \sin(z_1), \quad z_i \sim \mathcal{U}(-\pi, \pi)$$
$$f^{(2)} = \sin(z_1) + 4.75\sin^2(z_2) + \frac{1}{10} z_3^4 \sin(z_1),$$
$$f^{(3)} = \sin(z_1) + 3\sin^2(z_2) + \frac{9}{10} z_3^2 \sin(z_1).$$

For this model, expectation and variance can be computed analytically as $\mathrm{E}[f(z)] = 2.5$ and $\mathrm{Var}[f(z)] \approx 10.845$. The reference values of the Sobol indices can be found in [3]. We study numerically the errors in the proposed multifidelity estimators. For the expectation and variance estimators, we use both approaches of linear correlation and nonlinear statistical dependence. Moreover, we consider the Sobol indices as a vector-valued output, using the multidimensional estimator proposed in Section 3.1.

4.1.1. Preprocessing budget

Fig. 1 shows the histograms of the estimated optimal m with 5000 replicates of the preprocessing step. In each replicate, m is obtained by analytic minimization (7) of the error (15), computed with either 20 (blue) or 100 (red) samples. Fig. 2 shows the histogram of the computed estimators for 2000 replicates of the full analysis, again performed with either 20 (blue) or 100 (red) preprocessing samples. From Fig. 1, it is clear that the number of samples used for preprocessing has an impact on the accuracy of the computed m (and α). This is due to the accuracy of the variance and correlation estimates, which appear in the formula for computing m and α. Fortunately, the second figure shows that this does not have an impact on the computed estimators. In all of the examples, we have used 100 samples as preprocessing budget.

4.1.2. Expectation and variance estimation

Fig. 3 shows the computed error of the expectation and variance estimators using 100 replicates of the proposed estimators (13) and (16). We see that in this case, taking advantage of the nonlinear statistical dependence has no advantage over using only linear correlations. Table 1 summarizes the optimal α and m for the case of linear correlations, both averaged over 100 runs. As discussed in [3], the optimal budget allocations for expectation and variance are similar. Therefore, in principle there is no advantage in using a separate allocation for the variance. However, for this problem, we get a better approximation of the optimal α (0.8826 vs. 0.9455, as given in [3]). This improves the accuracy of the computed variance

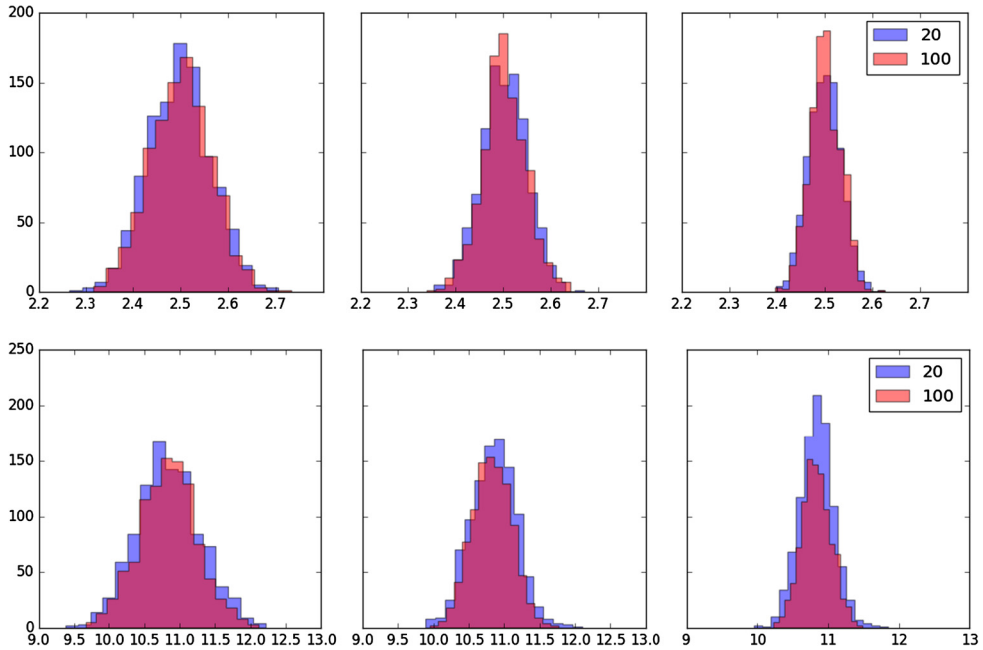


Fig. 2. Distribution of the computed estimator of the expectation (top row) and variance (bottom row) for the Ishigami model with 20 (blue) and 100 (red) preprocessing samples, computed with 1000 replicates. Columns: $p = 40$, $p = 80$, and $p = 160$. Plot axes: values (horizontal) and frequency (vertical) of the estimator.

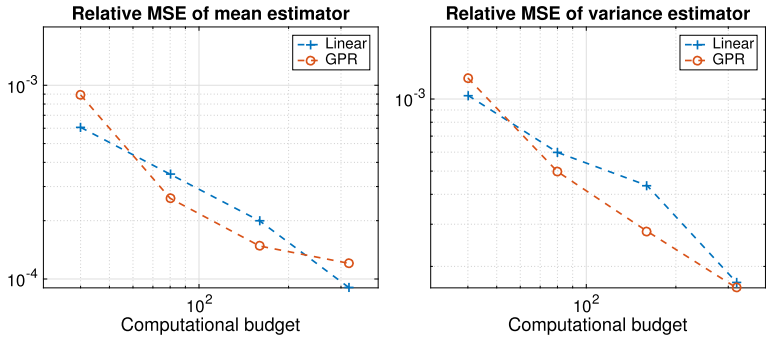


Fig. 3. Relative MSE of the proposed expectation and variance estimators for the Ishigami model as a function of computational budget p, computed with 100 replicates, for the linear and nonlinear estimators. Blue line: estimator using equation (13). Red line: estimator using equation (16) and a GPR model is trained with 100 samples.

Table 1
Estimates for the optimal α and m (with a budget $p = 40$) of the Ishigami model obtained via analytic minimization of (15) for the case of linear correlations. First two columns: values for the expectation estimator. Last two columns: values for the variance estimator. Both are averaged over 100 runs and estimated using 100 preprocessing samples.

Models	Expectation est.		Variance est.	
	m_k	α_k	m_k	α_k
$f^{(1)}$	7	1	8	1
$f^{(2)}$	461	1.0144	458	1.0144
$f^{(3)}$	9633	0.8826	9564	0.9289

with no extra cost, since the samples used for estimating ρ and σ in the standard multifidelity estimators, which are required to estimate the expectation, can be re-used for the proposed generalized ones. Moreover, this is achieved without using any a priori error estimate and can therefore be used also in the cases where such bounds are not available.

Table 2
Estimates for the optimal α and m (with a budget $p = 40$) obtained via analytic minimization of (15) for the main (columns 2-5) and total (columns 6-9) Sobol index estimators. All values are averaged over 100 runs and estimated using 100 preprocessing samples.

Models	Main index est.				Total index est.			
	m_k^M	α_k^M	α_k^M	α_k^M	m_k^T	α_k^T	α_k^T	α_k^T
$f^{(1)}$	9	1	1	1	0	–	–	–
$f^{(2)}$	485	1.049	1.057	1.028	55	1	1.108	1
$f^{(3)}$	7071	1.005	0.976	0.922	12471	0.828	2.778	1.051

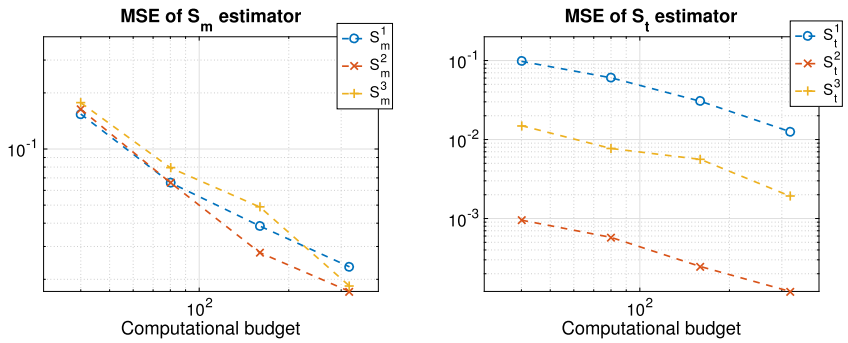


Fig. 4. Absolute MSE of the proposed Sobol estimators for the Ishigami model as a function of the computational budget p, computed with 100 replicates. Left: main effect sensitivities. Right: total effect sensitivities. Both are computed with the linear vector-valued estimator.

4.1.3. Sobol indices estimation

Here, we consider the main and total effect Sobol indices as vector-valued outputs of three components each. Therefore, we obtained a single optimal value m_k for each set of sensitivities, as discussed in Section 3.1, rather than a separate optimum for each index. Table 2 summarizes the estimates for the optimal α and m, both averaged over 100 runs. We see that the optimal allocation differs between the two sets, and also with respect to the expectation and variance, shown in Table 1. Therefore, in this case the optimum for the expectation does not guarantee a nearly optimal allocation of the samples. Fig. 4 shows the computed error of the proposed multifidelity Sobol indices estimators using 100 runs of the multifidelity algorithm, using only linear correlations. In particular, the errors are similar for the main effect sensitivities, while they differ for the total effect ones. We have used uniform weights $|\Omega_j| = 1$ in (10), however a different choice would have been possible if a higher precision had been required for one of the Sobol indices.

4.2. Quintic model

We here consider the following model

$$f^{(1)} = \sin(z_1) + \sin^2(z_2) + \frac{1}{10} z_3^5, \quad z_i \sim \mathcal{U}(-\pi, \pi)$$
$$f^{(2)} = \sin(z_1) + \sin^2(z_2) + 2z_3^3,$$
$$f^{(3)} = \sin(z_1) + \sin^2(z_2) + 20z_3.$$

For this model, we estimated the exact expectation and variance using a standard Monte Carlo method. For the proposed multifidelity estimators of the expectation and variance, we use both approaches of linear correlation and statistical dependence, which are referred to in Figs. 5 and 6 as linear and GPR, respectively. For both strategies, we also give the optimal budget allocations in Tables 3 and 4.

4.2.1. Expectation and variance estimation

Tables 3 and 4 summarize the values ρ and the optimal α and m for the case of, respectively, linear correlations and nonlinear statistical dependence. Both are averaged over 100 runs and the nonlinear model is estimated using 100 preprocessing samples. We see that in particular, the use of the nonlinear model significantly increases the correlation, to the point that the second model becomes redundant for estimating the expectation. Similarly, for the variance the number of model evaluations required by the high-fidelity reduces from 24 to 14 when the nonlinear dependence is exploited. Fig. 6 shows the computed error of the expectation and variance estimators as a function of the budget, using 100 replicates of the multifidelity algorithm. From this chart, it is evident that taking advantage of the nonlinear statistical dependence, as done by the estimator (16), provides a significant error reduction over the use of the more standard multifidelity estimator (13), which is only based on linear correlations.

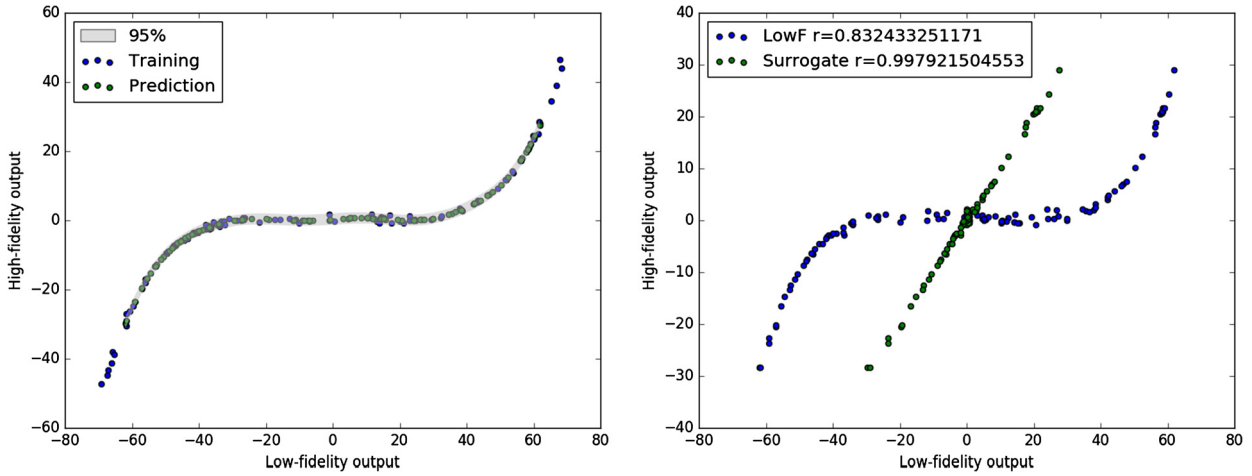


Fig. 5. Left: Training (blue) and validation (green) used to train the GPR model mapping between the low-fidelity $f^{(2)}$ and high-fidelity $f^{(1)}$ models for the expectation of the Quintic problem, with confidence intervals of the prediction (gray). Right: Comparison of resulting linear correlations when using a linear (blue) or a nonlinear (green) function for mapping between the low-fidelity $f^{(2)}$ and high-fidelity $f^{(1)}$ models.

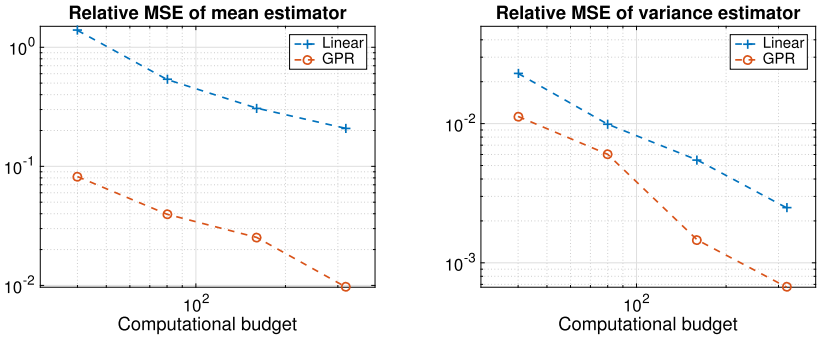


Fig. 6. MSE of the proposed expectation and variance estimators for the Quintic model as a function of the number of computational budget p, computed with 100 replicates, for the linear and nonlinear estimators. Blue line: estimator using equation (13). Red line: estimator using equation (16) and a GPR model is trained with 100 samples.

Table 3
Estimated optimal α and m (with a budget $p = 40$) of the Quintic model obtained via analytic minimization of (15) for the case of linear correlations. First three columns: values for the expectation estimator. Last three columns: values for the variance estimator. Both are averaged over 100 runs and estimated using 100 preprocessing samples.

Models	Expectation est.			Variance est.		
	m_k	α_k	ρ_k	m_k	α_k	ρ_k
$f^{(1)}$	25	1	1	24	1	1
$f^{(2)}$	254	0.384	0.973	263	0.195	0.972
$f^{(3)}$	2823	0.209	0.889	2521	0.119	0.784

5. Application to cardiac electrophysiology

We apply here the multifidelity methodology to a problem from cardiac electrophysiology. Briefly, the heart is an electrically active organ, with an electric wave spreading throughout the tissue thus to dictate the mechanical contraction. The activation pattern, defined as the first arrival time of the wave, is of clinical interest, as an abnormal pattern leads to a mechanical dysfunction. The propagation of the traveling wave is anisotropic, being faster along the cardiac fibers (usually 3- to 6-fold the transverse conduction velocity). Fiber distribution is roughly known from histological studies [20]. Patient-specific models incorporate the fiber distribution either from diffusion-tensor MRI [21] or tailored to the anatomy with a rule-based approach [22,23]. In both cases, epistemic uncertainty is present.

In what follows, we consider the problem of quantifying the uncertainty in the activation map, named $\psi(x)$, $x \in \Omega$ under the assumption of randomly distributed cardiac fibers, $\mathbf{F}(x)$.

Table 4
Estimated optimal α and m (with a budget $p = 40$) of the Quintic model obtained via analytic minimization of (18) for the case of nonlinear correlations. First three columns: values for the expectation estimator. Last three columns: values for the variance estimator. Both are averaged over 100 runs and estimated using 100 preprocessing samples.

Models	Expectation est.			Variance est.		
	m_k	α_k	ρ_k	m_k	α_k	ρ_k
$f^{(1)}$	28	1	1	14	1	1
$f^{(2)}$	0	–	–	417	0.971	0.996
$f^{(3)}$	11768	0.99	0.997	5469	0.846	0.865

The high-fidelity model is the monodomain equation [24]. This is a reaction-diffusion equation, with nonlinear reaction, coupled to a possibly large system of ODEs, encoding for the cellular membrane model. The solution of the problem are the spatio-temporal evolution of the transmembrane potential $V_{\mathrm{m}}(x,t)$ and some auxiliary membrane variables. The activation map is eventually obtained from the potential as the time of first crossing of a given threshold:

$$\psi_{\mathrm{h}}(x) = \inf_t \{t\colon V_{\mathrm{m}}(x,t) \geq V_{\mathrm{m,thres}}\}.$$

The monodomain equation reads as follows:

$$\beta\left(C_{\mathrm{m}}\frac{\partial V_{\mathrm{m}}}{\partial t} + I_{\mathrm{ion}}(V_{\mathrm{m}}, \mathbf{z}) - I_{\mathrm{stim}}(x,t)\right) = \nabla\cdot(\mathbf{G}_{\mathrm{m}}\nabla V_{\mathrm{m}}),$$

where C_{m} is the membrane capacitance, β the surface-to-volume ratio, $\mathbf{z}$ the set of auxiliary variables, and $\mathbf{G}_{\mathrm{m}}$ the monodomain conductivity tensor, namely

$$\mathbf{G}_{\mathrm{m}}(x) = \sigma_{\mathrm{t}}(x)\mathbf{I} + \big(\sigma_{\mathrm{l}}(x) - \sigma_{\mathrm{t}}(x)\big)\mathbf{F}(x)\otimes\mathbf{F}(x),$$

with σ_{l} and σ_{t} respectively the longitudinal and transverse conductivity.

The low-fidelity model is the zeroth-order eikonal approximation of the monodomain equation [25,24]. The eikonal model is obtained by assuming that $V_{\mathrm{m}}(x,t) = U(t - \psi_{\mathrm{l}}(x))$ for some action potential shape $U(\xi)$ and activation map ψ_{l}, and then neglecting the diffusion term. The model reads:

$$\begin{cases} \dfrac{\theta}{\sqrt{\beta}}\sqrt{\mathbf{G}_{\mathrm{m}}\nabla\psi\cdot\nabla\psi} = 1, \text{in } \Omega\setminus\{x_i\}_{i=1}^K, \\ \psi(x_i) = \tau_i, \qquad\qquad i = 1,\dots,K. \end{cases} \tag{20}$$

The boundary conditions set the initial activation time to τ_i at some given locations $x_i \in \Omega$. The parameter θ is a scaling factor linked to the membrane model. We remark that for the purposes of the analysis the function $U(\xi)$ is irrelevant.

In this study, we consider a patient-specific anatomy. The electrophysiology parameters were already fitted to the patient data in a previous study [26], with good correlation in terms of activation map and surface ECG. The numerical grid is composed by approximately 20 millions voxels at full resolution (0.2 mm), and 200 000 voxels at coarse resolution (1 mm).

For the monodomain simulation, we used Propag-5 software [27]. Only the full resolution grid is considered, necessary to ensure small numerical error [28,29]. The membrane model is from ten Tusscher et al. [30]. Activation is initiated with a current stimulus in a 1 mm^3 region for 2 ms. Regarding the eikonal model, we used our GPGPU implementation described in [25]. We performed eikonal simulations at both coarse and full resolution.

In summary, the full hierarchy of models for the multifidelity framework includes: the discretized monodomain model, $f^{(1)}\colon \mathbf{F}\mapsto\psi_{\mathrm{h}}$, the discretized eikonal model with high spatial resolution, $f^{(2)}\colon \mathbf{F}\mapsto\psi_{\mathrm{l},1}$, and the eikonal model at coarse spatial resolution, $f^{(3)}\colon \mathbf{F}\mapsto\psi_{\mathrm{l},2}$.

5.1. Random fiber distribution and sampling

The fiber distribution is modeled through a rule-based approach. In details, the fiber orientation is locally determined by 3 angles, namely α, ϕ and γ. These angles are related to the Euler's angles, but conveniently adapted so that α is consistent with the definition given by Streeter et al. [20]: it is the angle between the local fiber direction and the valves' plane of the heart. Transmurally, the angle α varies cubically from $\pi/3$ at the endocardium to $-\pi/3$ at the epicardium [31]

$$\alpha(r) = \frac{\pi}{3}(1-2r)^3, \tag{21}$$

where r measures the endocardium-to-epicardium distance, normalized to unity. ($r = 0$ corresponds to the endocardium, $r = 1$ to the epicardium.)

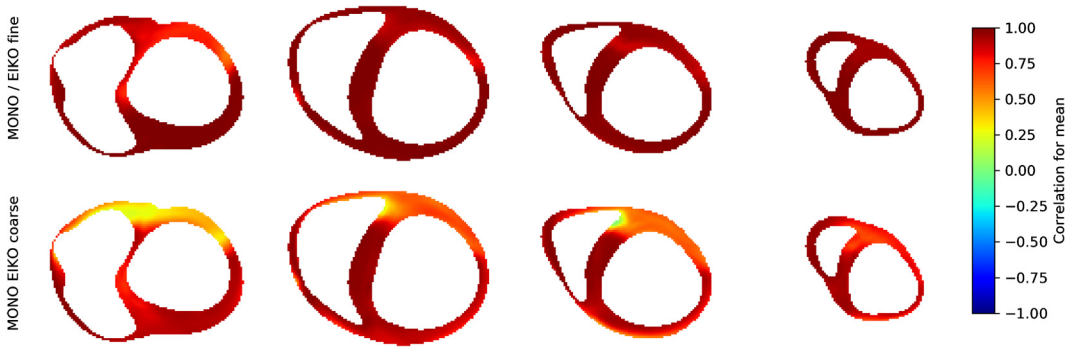


Fig. 7. Correlations for the expectation estimator between monodomain and eikonal models at the same mesh resolution $h = 0.02$ cm (top row) and at a coarser resolution $h = 0.1$ cm (bottom row). Correlations are reported for slices of the heart from basal (left) to apical plane (right).

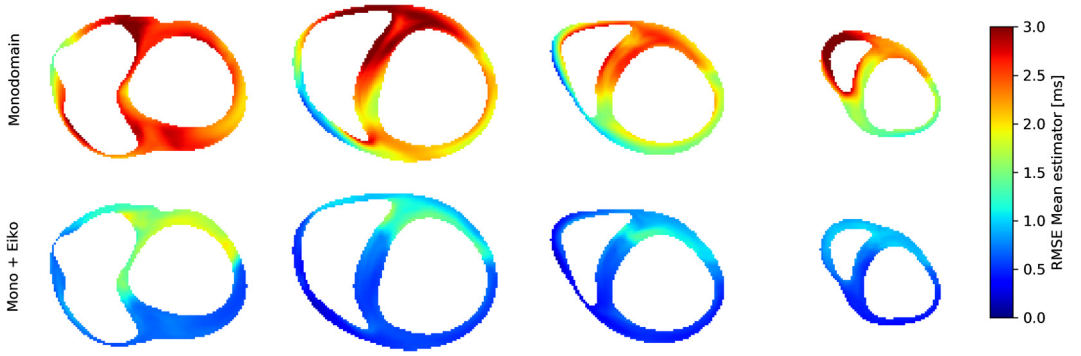


Fig. 8. Estimated RMSE of the expectation estimator when only the monodomain model is used (top row), and its reduction when the eikonal model with fine grid is added (bottom row). Errors are reported for slices of the heart from basal plane (left) to apical plane (right).

The random fiber field is obtained with the above rule, but assuming that the angle α is a Gaussian random (scalar field), named $A(x)$, with mean given by (21) and correlation function as follows

$$k(x, y) = \sigma \exp\Big(-\frac{d(x, y)^2}{2\ell^2}\Big),$$

where σ is the standard deviation, ℓ the correlation length, and $d(x, y)$ the geodesic distance between x and y.

The sampling strategy is provided in details in [17]. In summary, the random angle $A(x)$ is approximated by a truncated Karhuhen-Loève expansion:

$$A(x) = A_0(x) + \sigma \sum_{i=1}^{K} \sqrt{\lambda_i} Z_i \varphi_i(x),$$

with Z_i independent standard normal random variables, and λ_i and φ_i being eigenvalues and eigenvectors of the Hilbert-Schmidt operator associated to the kernel $k(x, y)$. The optimal truncation is obtained from a low-rank approximation of the discretized correlation matrix via a pivoted Cholesky decomposition [32]. The distance function, required during the assembly of the correlation matrix, is evaluated by solving an eikonal problem.

5.2. Expectation estimation

Fig. 7 shows the correlations between the expectation estimators for the bidomain and eikonal models at the same mesh resolution $h = 0.02$ cm (top) and at a coarser resolution $h = 0.1$ cm (bottom). Fig. 8 shows the estimated error in the estimated expectation when only the bidomain model is used (top) and its reduction when the fine eikonal model is added (bottom). Fig. 11 (top) shows the computed expectation when all three models are used for the estimation. Table 5 summarizes the optimal number of model evaluations, estimated RMSEs e_k with increasing number of models, and correlations $\bar{\rho}_k$ for a tolerance of $\epsilon = 0.75$ ms. As a comparison, standard MC would require 92 samples of the high-fidelity model in order to achieve the same accuracy. The estimated node time for standard and multifidelity MC is, respectively, 277 h and 36 h.

Table 5
Estimated optimal number of model evaluations m, RMSEs e_k of the electrophysiology model obtained via analytic minimization of (11) with increasing number of models, and correlations $\bar{\rho}_k$. First two columns: values for the expectation estimator (tolerance of $\epsilon = 0.75$ ms). Last two columns: values for the variance estimator (tolerance of $\epsilon = 10.5$ ms^2).

Models	Expectation est.			Variance est.		
	m_k	e_k	$\bar{\rho}_k$	m_k	e_k	$\bar{\rho}_k$
$f^{(1)}$	10	2.28 ms	1	10	20.34 ms^2	1
$f^{(2)}$	165	0.87 ms	0.953	140	9.06 ms^2	0.905
$f^{(3)}$	2490	0.76 ms	0.750	1090	8.07 ms^2	0.567

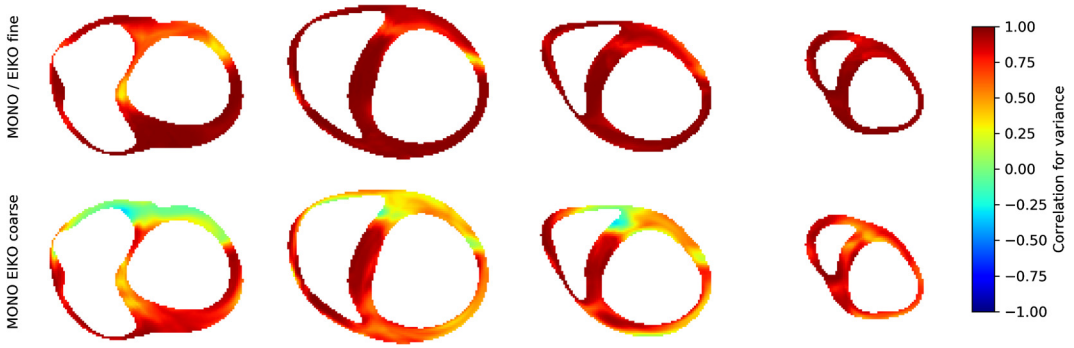


Fig. 9. Correlations for the variance estimator between monodomain and eikonal models at the same mesh resolution $h = 0.02$ cm (top row) and at a coarser resolution $h = 0.1$ cm (bottom row). Correlations are reported for slices of the heart from basal (left) to apical plane (right).

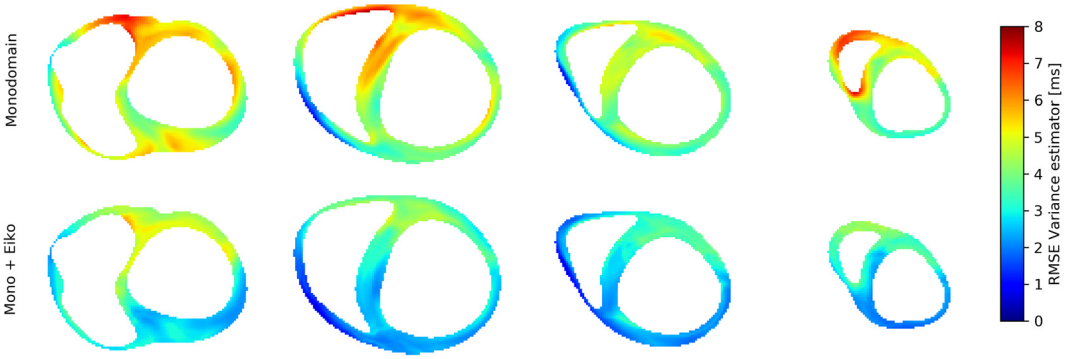


Fig. 10. Estimated RMSE of the variance estimator when only the monodomal model is used (top row), and its reduction when the eikonal model with fine grid is added (bottom row). Errors are reported for slices of the heart from basal plane (left) to apical plane (right).

5.3. Variance estimation

Fig. 9 shows the correlations between the variance estimators for the bidomain and eikonal models at the same mesh resolution $h = 0.02$ cm (top) and at a coarser resolution $h = 0.1$ cm (bottom). Fig. 10 shows the estimated error in the estimated standard deviation when only the bidomain model is used (top) and its reduction when the fine eikonal model is added (bottom). Fig. 11 (bottom) shows the computed standard deviation when all three models are used for the estimation. Table 5 summarizes the optimal number of model evaluations, estimated RMSEs e_k with increasing number of models, and correlations $\bar{\rho}_k$ for a tolerance of $\epsilon = 10.5$ ms^2. As a comparison, standard MC would require 48 samples of the high-fidelity model in order to achieve the same accuracy. The estimated node time for standard and multifidelity MC is, respectively, 145 h and 35 h.

6. Application to nonlinear mechanics

The last application is inspired by to the *indentation problem* of hyperelastic material, as described for instance in [REF]. Briefly, we consider a $2 \times 2 \times 1\,\text{mm}^3$ cube clamped at the bottom, while top face is indented of 0.1 mm in a $1\,\text{mm}^2$ square region at its center. We consider the finite deformation regime of a incompressible neo-Hookean material with strain energy density function given as follows:

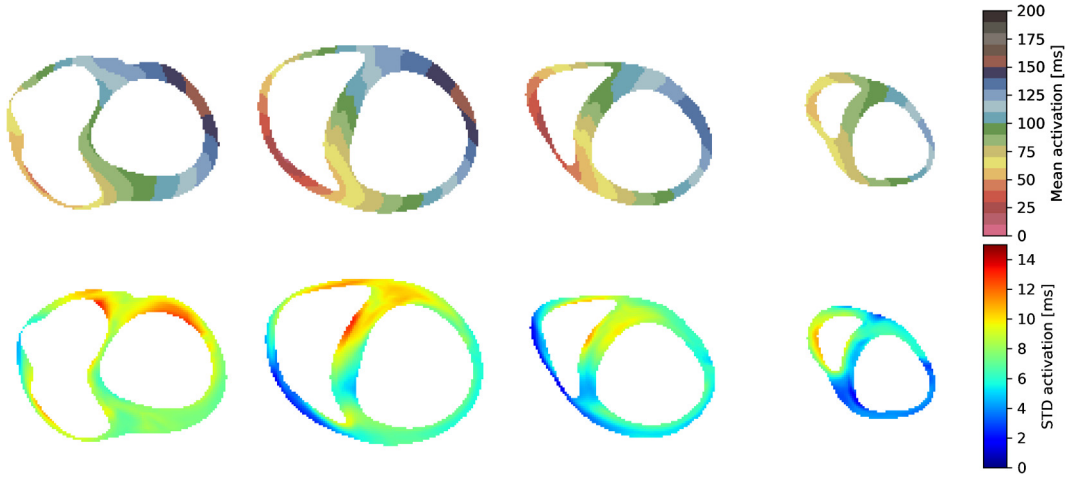


Fig. 11. Estimated average of the activation map (top row) and estimated standard deviation (bottom row) when all three models are used.

$$\mathcal{W}(\mathbf{F}) = \frac{\mu}{2}(\mathbf{F}:\mathbf{F} - 3) - \mu \ln(J).$$

The deformation gradient tensor is defined as $\mathbf{F} = \mathbf{I} + \nabla \mathbf{u}$, where $\mathbf{u}\colon \Omega \to \mathbb{R}^3$ is the (unknown) displacement. The shear modulus is set to vary in space following a spatially-correlated random field, evaluated via Karhuhen-Loève expansion:

$$\mu(x) = \mu_0 + \sigma \sum_{i=1}^{K} \sqrt{\lambda_i} Z_i \phi_i(x), \qquad Z_i \sim \mathcal{N}(0,1).$$

Functions $\phi_i(x)$ and coefficients λ_i are, respectively, eigenfunctions and eigenvalues of the Hilbert-Schmidt operator for the squared-exponential correlation function, as in the previous section. The correlation length is $\ell = \frac{1}{2\sqrt{2}}$, while mean and standard deviation of μ are set to 80 MPa and 20 MPa. Samples with non-positive $\mu(x)$ are discarded.

The aim of this experiment is to estimate the pointwise average of the von Mises stress, which is related to the second invariant of the Cauchy stress:

$$S = \frac{\mu}{J}\sqrt{\tfrac{3}{2}\,\mathrm{dev}\,\mathbf{B} : \mathrm{dev}\,\mathbf{B}}, \qquad \mathbf{B} = \mathbf{F}\mathbf{F}^{\mathrm{T}},$$

given the uncertainty in the shear modulus.

The high-fidelity model is the Finite Element (FE) discretization of the problem in mixed form with Taylor-Hood $\mathbb{P}^2 - \mathbb{P}^1$ elements. We introduce the pressure field p to enforce the incompressibility constraint $J = 1$. The discrete problem is then solved with the Newton's method and a direct solver for the linearized system.

The low-fidelity problem is based on the *compressible* version of the neo-Hookean material,

$$\mathcal{W}_{\mathrm{lf}}(\mathbf{F}) = \frac{\mu}{2}(\mathbf{F}:\mathbf{F} - 3) - \mu \ln(J) + \frac{\lambda}{2}(J-1)^2,$$

with $\lambda = 10\mu_0 = 800$ MPa. We discretized the problem with linear FE for the displacement. Since the Poisson's ratio is about 0.45, we do not expect locking for this setup. In addition, the linear system within each Newton's iteration is solved via algebraic multigrid.

6.1. Expectation estimation of Von Mises stresses

Fig. 12 (left) shows the computed expectation when both the high- and low-fidelity models are used for the estimation. Fig. 12 (center) shows the correlations between the expectation estimators for the two models. Fig. 12 (right) shows the estimated error in the estimated expectation when both models are used. We note that the high concentration of the error along the displaced boundary would have significantly increased the number of required model evaluations, had each point been treated as a separate scalar-valued QoI. Table 6 summarizes the optimal number of model evaluations, estimated RMSEs e_k with increasing number of models, and correlations $\bar{\rho}_k$ for a budget of $B = 2000$ s. As a comparison, standard MC would require 293 samples of the high-fidelity model in order to achieve the same accuracy. The estimated CPU time for standard and multifidelity MC is, respectively, 27000 s and 2000 s.

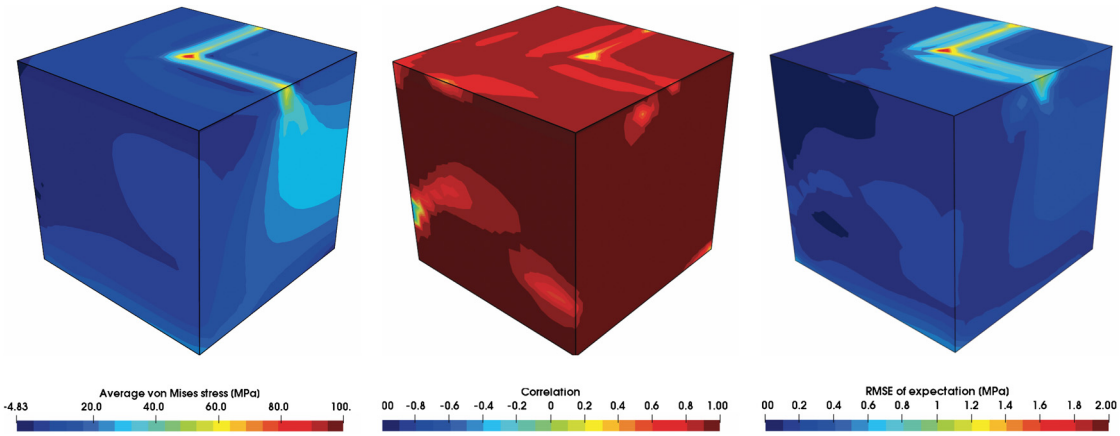


Fig. 12. MFMC estimate of average von Mises stress (left), pointwise correlation coefficient (center), and RMSE of the estimator (right).

Table 6
Estimated optimal number of model evaluations m, RMSEs e_k of the nonlinear mechanics model obtained via analytic minimization of (11) with increasing number of models, and correlations $\bar{\rho}_k$. First three columns: values for the expectation estimator (budget $B = 2000$ s, $w_1 = 86$ s, $w_2 = 1.85$ s). Last three columns: values for the variance estimator (same budget as previous set).

Models	Expectation est.			Variance est.		
	m_k	e_k	$\bar{\rho}_k$	m_k	e_k	$\bar{\rho}_k$
$f^{(1)}$	10	1.23 MPa	1	10	12.26 MPa^2	1
$f^{(2)}$	550	0.223 MPa	0.992	395	2.85 MPa^2	0.985

6.2. Variance estimation of Von Mises stresses

Table 6 (last three columns) summarizes the optimal number of model evaluations, estimated RMSEs e_k with increasing number of models, and correlations $\bar{\rho}_k$ for a budget of $B = 2000$ s. As a comparison, standard MC would require 184 samples of the high-fidelity model in order to achieve the same accuracy. The estimated CPU time for standard and multifidelity MC is, respectively, 16000 s and 2000 s. The qualitative profiles of the estimated variance, the correlation, and the error, follow closely those of the expectation and are therefore not reported in the figures.

7. Conclusions

In this work, we extended and generalized previously proposed multifidelity Monte Carlo estimators for the expectation [2], the variance, and sensitivity indices [3]. We tested our methodology on selected benchmarks, showing the good performance and usability of the proposed estimators. Moreover, we detailed their application to a real-world example from cardiac electrophysiology, which is unapproachable by standard Monte Carlo techniques.

Rather than considering the low- and high-fidelity models as opposing, which should be selected according to time constraints or accuracy considerations, the multifidelity framework combines them and exploits their best features. By doing so, a full UQ analysis of the complex problems, such as the monodomain equation in cardiac electrophysiology, can be performed in a total time comparable to that of using only the low-fidelity model, without sacrificing the complexity of the high-fidelity solution.

Future works will focus on reduced-physics models for electrophysiology, incorporating the repolarization and the ECG into the multifidelity hierarchy. Moreover, we intend to study multifidelity estimators for inverse problems.

Acknowledgements

The authors acknowledge financial support by the Theo Rossi di Montelera Foundation, the Metis Foundation Sergio Mantegazza, the Fidinam Foundation, and the Horten Foundation to the Center for Computational Medicine in Cardiology. This work was also supported by grants from the Swiss National Supercomputing Centre (CSCS) under project ID s778.

References

[1] B. Peherstorfer, K. Willcox, M. Gunzburger, Survey of multifidelity methods in uncertainty propagation, inference, and optimization. ArXiv e-prints, June 2018.

[2] Benjamin Peherstorfer, Karen Willcox, Max Gunzburger, Optimal model management for multifidelity monte carlo estimation, SIAM J. Sci. Comput. 38 (5) (January 2016) A3163–A3194, https://doi.org/10.1137/15m1046472.
[3] E. Qian, Benjamin Peherstorfer, D. O'Malley, V.V. Vesselinov, K. Willcox, Multifidelity monte carlo estimation of variance and sensitivity indices, SIAM/ASA J. Uncertain. Quantificat. 6 (2) (2018) 683–706, https://doi.org/10.1137/17M1151006.
[4] Steven Niederer, Lawrence Mitchell, Nicolas Smith, Gernot Plank, Simulating human cardiac electrophysiology on clinical time-scales, Front. Physiol. 2 (2011).
[5] Pras Pathmanathan, Richard A. Gray, Vicente J. Romero, Tina M. Morrison, Applicability analysis of validation evidence for biomedical computational models, J. Verif. Valid. Uncertain. Quantificat. 2 (2) (September 2017) 021005, https://doi.org/10.1115/1.4037671.
[6] Michael B. Giles, Multilevel monte carlo methods, Acta Numer. 24 (April 2015) 259–328, https://doi.org/10.1017/S096249291500001X.
[7] George Fishman, Monte Carlo: Concepts, Algorithms, and Applications, corrected edition, Springer, New York, ISBN 978-0-387-94527-9, February 2003.
[8] Jonas Biehler, Michael W. Gee, Wolfgang A. Wall, Towards efficient uncertainty quantification in complex and large-scale biomechanical problems based on a bayesian multi-fidelity scheme, Biomech. Model. Mechanobiol. 14 (3) (2015) 489–513.
[9] Youssef Rouchdy, Jérôme Pousin, Joël Schaerer, Patrick Clarysse, A nonlinear elastic deformable template for soft structure segmentation: application to the heart segmentation in MRI, Inverse Probl. 23 (3) (April 2007) 1017–1035, https://doi.org/10.1088/0266-5611/23/3/011.
[10] Michele Pisaroni, Sebastian Krumscheid, Fabio Nobile, Quantifying Uncertain System Outputs via the Multilevel Monte Carlo Method–Part I: Central Moment Estimation, Technical report, MATHICSE Technical Report 23.2017, EPFL, 2017.
[11] Claudio Bierig, Alexey Chernov, Convergence analysis of multilevel monte carlo variance estimators and application for random obstacle problems, Numer. Math. 130 (4) (2015) 579–613.
[12] Jan S. Hesthaven, Gianluigi Rozza, Benjamin Stamm, et al., Certified Reduced Basis Methods for Parametrized Partial Differential Equations, Springer, 2016.
[13] Cesare Corrado, Jean-Frédéric Gerbeau, Philippe Moireau, Identification of weakly coupled multiphysics problems. Application to the inverse problem of electrocardiography, J. Comput. Phys. 283 (February 2015) 271–298, https://doi.org/10.1016/j.jcp.2014.11.041.
[14] Andrew J. Pullan, Karl A. Tomlinson, Peter J. Hunter, A finite element method for an eikonal equation model of myocardial excitation wavefront propagation, SIAM J. Appl. Math. 63 (1) (2002) 324–350, https://doi.org/10.1137/S0036139901389513.
[15] Piero Colli Franzone, Luciano Guerri, Spreading of excitation in 3-D models of the anisotropic cardiac tissue. I. Validation of the eikonal model, Math. Biosci. 113 (1993) 145–209, https://doi.org/10.1016/0025-5564(93)90001-q.
[16] Phaedon-Stelios Koutsourelakis, Accurate uncertainty quantification using inaccurate computational models, SIAM J. Sci. Comput. 31 (5) (January 2009) 3274–3300, https://doi.org/10.1137/080733565.
[17] A. Quaglino, S. Pezzuto, P.S. Koutsourelakis, A. Auricchio, R. Krause, Fast uncertainty quantification of activation sequences in patient-specific cardiac electrophysiology meeting clinical time constraints, Int. J. Numer. Methods Biomed. Eng. 34 (7) (2018) e2985, https://doi.org/10.1002/cnm.2985.
[18] Malvin H. Kalos, Paula A. Whitlock, Monte Carlo Methods, John Wiley & Sons, ISBN 978-3-527-40760-6, October 2008, Google-Books-ID: b8Xb4rBkygUC.
[19] Tsutomu Ishigami, Toshimitsu Homma, An importance quantification technique in uncertainty analysis for computer models, in: [1990] Proceedings. First International Symposium on Uncertainty Modeling and Analysis, IEEE Comput. Soc. Press, 1990, pp. 398–403.
[20] Daniel D. Streeter, Henry M. Spotnitz, Dali P. Patel, John Ross, Edmund H. Sonnenblick, Fiber orientation in the canine left ventricle during diastole and systole, Circ. Res. 24 (3) (1969) 339–347, https://doi.org/10.1161/01.res.24.3.339.
[21] Nicolas Toussaint, Christian T. Stoeck, Tobias Schaeffter, Sebastian Kozerke, Maxime Sermesant, Philip G. Batchelor, In vivo human cardiac fibre architecture estimation using shape-based diffusion tensor processing, Med. Image Anal. 17 (8) (2013) 1243–1255, https://doi.org/10.1016/j.media.2013.02.008.
[22] Mark Potse, Bruno Dubé, Jacques Richer, Alain Vinet, Ramesh M. Gulrajani, A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart, IEEE Trans. Biomed. Eng. 53 (12) (2006) 2425–2435, https://doi.org/10.1109/tbme.2006.880875.
[23] J.D. Bayer, R.C. Blake, G. Plank, N.A. Trayanova, A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models, Ann. Biomed. Eng. 40 (10) (2012) 2243–2254, https://doi.org/10.1007/s10439-012-0593-5.
[24] Piero Colli Franzone, Luca F. Pavarino, Simone Scacchi, Mathematical Cardiac Electrophysiology, Springer, 2014.
[25] Simone Pezzuto, Peter Kal'avský, Mark Potse, Frits W. Prinzen, Angelo Auricchio, Rolf Krause, Evaluation of a rapid anisotropic model for ecg simulation, Front. Physiol. (ISSN 1664-042X) 8 (2017) 265, https://doi.org/10.3389/fphys.2017.00265.
[26] Mark Potse, Dorian Krause, Wilco Kroon, Romina Murzilli, Stefano Muzzarelli, François Regoli, Enrico Caiani, Frits W. Prinzen, Rolf Krause, Angelo Auricchio, Patient-specific modelling of cardiac electrophysiology in heart-failure patients, Europace 16 (2014) iv56–iv61, https://doi.org/10.1093/europace/euu257.
[27] Thomas Dickopf, Dorian Krause, Rolf Krause, Mark Potse, Design and analysis of a lightweight parallel adaptive scheme for the solution of the monodomain equation, SIAM J. Sci. Comput. 36 (2) (2014) C163–C189, https://doi.org/10.1137/130912505.
[28] Simone Pezzuto, Johan Hake, Joakim Sundnes, Space-discretization error analysis and stabilization schemes for conduction velocity in cardiac electrophysiology, Int. J. Numer. Methods Biomed. Eng. (2016), https://doi.org/10.1002/cnm.2762.
[29] Pras Pathmanathan, Richard A. Gray, Verification of computational models of cardiac electro-physiology, Int. J. Numer. Methods Biomed. Eng. 30 (5) (November 2014) 525–544, https://doi.org/10.1002/cnm.2615.
[30] K.H.W.J. ten Tusscher, D. Noble, P.J. Noble, A.V. Panfilov, A model for human ventricular tissue, Am. J. Physiol., Heart Circ. Physiol. 286 (4) (2004) H1573–H1589, https://doi.org/10.1152/ajpheart.00794.2003.
[31] R. Beyar, S. Sideman, A computer study of the left ventricular performance based on fiber structure, sarcomere dynamics, and transmural electrical propagation velocity, Circ. Res. 55 (3) (1984) 358–375, https://doi.org/10.1161/01.res.55.3.358.
[32] Helmut Harbrecht, Michael Peters, Reinhold Schneider, On the low-rank approximation by the pivoted cholesky decomposition, Appl. Numer. Math. 62 (4) (2012) 428–440, https://doi.org/10.1016/j.apnum.2011.10.001.