Improving Multi-Model Trajectory Simulation Estimators using Model Selection and Tuning

Geoffrey F. Bomarito[§]
NASA Langley Research Center, Hampton, VA 23681, USA

Gianluca Geraci*
Sandia National Laboratories, Albuquerque, NM 87185, USA

James E. Warner[†], Patrick E. Leser[§], and William P. Leser[§]
NASA Langley Research Center, Hampton, VA 23681, USA

Michael S. Eldred[‡] and John D. Jakeman[§]
Sandia National Laboratories, Albuquerque, NM 87185, USA

Alex A. Gorodetsky[¶]
University of Michigan, Ann Arbor, MI 48109, USA

Multi-model Monte Carlo methods have been demonstrated to be an efficient and accurate alternative to standard Monte Carlo (MC) in the model-based propagation of uncertainty in entry, descent, and landing (EDL) applications. These multi-model MC methods fuse predictions from low-fidelity models with the high-fidelity EDL model of interest to produce unbiased statistics with a fraction of the computational cost. The accuracy and efficiency of the multi-model MC methods are dependent upon the magnitude of correlations of the low-fidelity models with the high-fidelity model, but also upon the correlation among the low-fidelity models, and their relative computational cost. Because of this layer of complexity, the question of how to optimally select the set of low-fidelity models has remained open. In this work, methods for optimal model construction and tuning are investigated as a means to increase the speed and precision of trajectory simulation for EDL. Specifically, the focus is on the inclusion of low-fidelity model tuning within the sample allocation optimization that accompanies multi-model MC methods. Preliminary results indicate that low-fidelity model tuning can significantly improve efficiency and precision of trajectory simulations and provide an increased edge to multi-model MC methods when compared to standard MC. The challenges and potential benefits to exploring a fully iterative and comprehensive optimization strategy in future work are highlighted.

I. Introduction

The National Aeronautics and Space Administration's (NASA) current concepts for human-class missions to Mars require the precision landing of multiple payloads; these payloads must be placed within 50m of their targets to balance surface travel required by humans and separation from each other to avoid damage. The successful Mars 2020 mission, equipped with the terrain-relative navigation technology, achieved a landing precision of 1-2 km of the preflight target, and within meters of the autonomously chosen in-flight target. Both of these are tremendous achievements and establish a new state of the art. However, for reaching preflight targets, new technologies in entry, descent, and landing (EDL) are needed to make up the two orders of magnitude improvement required for future precision landing [1]. The current

[§]Materials Engineer, Durability, Damage Tolerance, and Reliability Branch, M/S 188E.
*Senior Member of Technical Staff, Optimization and Uncertainty Quantification, Sandia National Laboratories, MS1318, Albuquerque, NM 87185, AIAA Member.
[†]Research Computer Scientist, Durability, Damage Tolerance, and Reliability Branch, M/S 188E.
[‡]Distinguished Member of Technical Staff, Optimization and Uncertainty Quantification, Sandia National Laboratories, MS1318, Albuquerque, NM 87185, AIAA Associate Fellow.
[§]Principal Member of Technical Staff, Optimization and Uncertainty Quantification, Sandia National Laboratories, Albuquerque, NM.
[¶]Assistant Professor, Department of Aerospace Engineering, University of Michigan, AIAA Member.

work investigates improvements to the accuracy and efficiency of trajectory simulation as a means of addressing this precision landing problem.

For the cases considered in this work, trajectory simulation consists of modeling the transit of a vehicle through an atmosphere to a landing location. The atmospheric conditions and initial angle of attack (among many other factors) control the quantities of interest (QOIs) of the trajectory simulation. Because these factors contain a significant amount of uncertainty, accurately predicting QOIs becomes a problem of uncertainty propagation. Monte Carlo (MC) dispersion analysis, wherein samples of the random factors are used to simulate several possible trajectories, is the standard approach for trajectory simulation [1, 2]. Because running high-fidelity EDL trajectory simulations can be time consuming, the use of MC can become cost prohibitive.

Recently it was shown that multi-model MC has the potential to significantly improve precision and speed of trajectory simulation [3]. In multi-model MC, lower-fidelity models are used in conjunction with a high-fidelity model to produce statistics of a QOI at a fraction of the cost. The key to multi-model MC is that the lower-fidelity models are correlated with the high-fidelity model and that statistical estimators can leverage the correlation to predict statistics of the QOI with more precision. Multi-level MC (MLMC) [4] and Multi-fidelity MC (MFMC) [5] are two of the most well-known multi-model MC methods. Both MLMC and MFMC make assumptions about the structure of the underlying models in their pursuit of optimal statistical estimators, namely they assume a hierarchical relationship among models. The generalized approximate control variate (ACV) framework was introduced to unify MLMC and MFMC into a single framework that permits more diverse structure/relationships among the high- and low-fidelity models [6]. It was shown that significant performance improvements could be obtained relative to MLMC and MFMC by optimizing for resource allocation within the more general ACV context.

Low-fidelity model development has been shown to be a key factor in the application of multi-model MC methods. Warner et al. [3] illustrated that the accuracy of a low-fidelity trajectory model was directly related to the amount of improvement seen when using multi-model MC trajectory simulation. Here, the performance of the approach failed to outperform standard MC for some dynamic EDL QOIs such as angle of attack along the trajectory because the low-fidelity models struggled to accurately predict them. Bomarito et al. [7] also illustrated that model selection can have a significant impact on performance of ACV estimators, including the optimal structure of the models and the accuracy of the resulting estimator. The underlying reason for this is that the accuracy and efficiency of ACVs are dependent upon the magnitude of correlations of the low-fidelity models with the high-fidelity model and also upon the correlation among the low-fidelity models. Especially in cases where models can be tuned, by controlling continuous or discrete hyperparameters, it is important to acknowledge the effect this may have on the ultimate performance of ACVs. With this motivation, the current work incorporates tunable low-fidelity models into ACV optimization.

The current work again revisits the analysis of the Adaptable Deployable Entry and Placement Technology (ADEPT) test flight [2, 3] to investigate how the construction and tuning of low-fidelity trajectory models impacts estimates of EDL QOIs. Program to Optimize Simulated Trajectories II (POST2), a legacy code developed at NASA Langley Research Center, is used for high-fidelity trajectory simulation. A prototype algorithm for the optimization of an ACV estimator is designed to provide initial feedback for an in-progress implementation of these capabilities within the Sandia uncertainty quantification (UQ) flagship code Dakota [8, 9]. For this initial demonstration, we consider two possible low-fidelity models, namely a model with reduced physics and/or a tunable-fidelity version of the POST2 simulator by reduced time integration.

The remainder of the paper is organized as it follows. In Section II we describe the main building block of our UQ approach: the ACV framework for the construction of generalized multi-model estimators. Sections III and IV introduce the possibility of selecting the model graph or the low-fidelity hyperparameters to achieve greater variance reductions. Numerical demonstrations for the EDL analyses are provided in Section V, where model tuning and graph adaptation are discussed and preliminary results are provided. Concluding remarks and perspective close the paper in Section VI.

II. Approximate Control Variates for Multi-Model Uncertainty Propagation

In this section we briefly introduce ACV estimators, which constitute the building block of our multi-model uncertainty propagation strategy. The standard MC algorithm is the most popular sampling method due to its simplicity, flexibility, and provably convergent behavior. If we denote a vector-valued random variable $Z \subset \mathbb{R}^d$, for a generic QOI $Q : \mathbb{R}^d \to \mathbb{R}$ the expected value can be approximated by collecting N realizations:

$$\mathbb{E}[Q] \simeq \hat{Q} = \frac{1}{N} \sum_{i=1}^{N} Q^{(i)}. \tag{1}$$

It is well-known that the MC estimator $\hat{Q}$ is itself a random variable, with its own mean and variance. While its mean is simply $\mathbb{E}[Q]$, since the estimator is unbiased, its variance depends on the QOI variance and the total number of evaluations N

$$\mathbb{V}ar\left[\hat{Q}\right] = \frac{\mathbb{V}ar[Q]}{N}, \tag{2}$$

which illustrates the difficulty of obtaining reliable statistics for high-fidelity, expensive computer codes, due to the limited number of realizations N that are usually available in this scenario.

ACV is a generalization of the optimal control variate (OCV) approach [10], where M low-fidelity models are introduced to reduce the estimator variance. The OCV estimator is obtained by adding an unbiased term $(\hat{Q}_i - \mu_i)$, where $\hat{Q}_i$ indicates the ith low-fidelity mean estimator and μ_i its known expected value. Each of these M terms is then weighted by coefficients $\underline{\alpha} = [\alpha_i, \dots, \alpha_M]^{\mathrm{T}}$ that are subsequently optimized to obtain a minimum variance for a prescribed cost:

$$\hat{Q}^{\mathrm{CV}} = \hat{Q} + \sum_{i=1}^{M} \alpha_i \left(\hat{Q}_i - \mu_i\right). \tag{3}$$

In this framework, the minimum variance is obtained by choosing $\underline{\alpha}^\star = -\mathbf{C}^{-1}\mathbf{c}$, where $\mathbf{C} \in \mathbb{R}^{M\times M}$ is the covariance matrix among the low-fidelity evaluations Q_i and $\mathbf{c} \in \mathbb{R}^M$ is the vector of covariances between the high-fidelity realizations Q and the Q_i low-fidelity realizations. The choice of optimal coefficients $\underline{\alpha}^\star$ corresponds to a variance reduction, with respect to MC, equal to

$$\gamma^{\mathrm{CV}}(\underline{\alpha}^\star) = \frac{\mathbb{V}ar\left[\hat{Q}^{\mathrm{CV}}\right]}{\mathbb{V}ar\left[\hat{Q}\right]} = 1 - R^2_{OCV} = 1 - \frac{\mathbf{c}^{\mathrm{T}}\mathbf{C}^{-1}\mathbf{c}}{\mathbb{V}ar[Q]}, \quad 0 \le R^2_{OCV} \le 1. \tag{4}$$

The ACV framework relaxes the hypothesis of knowing the expected values for each of the low-fidelity models, which is the most common situation in many engineering analysis scenarios. The approximation of μ_i is accomplished by partitioning the set of samples available for each model. If we consider $\mathbf{z}$ as the set of samples for the high-fidelity model and $\mathbf{z}_i$ is an ordered set of $N_i = \lceil r_i N \rceil$ evaluations for the ith low-fidelity model Q_i, we can partition $\mathbf{z}_i$ in two ordered subsets. More precisely, we can define two sets $\mathbf{z}_i^1$ and $\mathbf{z}_i^2$ with $\mathbf{z}_i^1 \bigcup \mathbf{z}_i^2 = \mathbf{z}_i$.

The general ACV estimator can be written as

$$\tilde{Q}\left(\underline{\alpha}, \underline{\mathbf{z}}\right) = \hat{Q}(\mathbf{z}) + \sum_{i=1}^{M} \alpha_i \left(\hat{Q}_i(\mathbf{z}_i^1) - \hat{\mu}_i(\mathbf{z}_i^2)\right) = \hat{Q}(\mathbf{z}) + \sum_{i=1}^{M} \alpha_i \Delta_i(\mathbf{z}_i) = \hat{Q} + \underline{\alpha}^{\mathrm{T}}\underline{\Delta}, \tag{5}$$

where $\underline{\Delta} = \left[\Delta_1(\mathbf{z}_1), \dots, \Delta_M(\mathbf{z}_M)\right]^{\mathrm{T}}$, $\Delta_i(\mathbf{z}_i) = \hat{Q}_i(\mathbf{z}_i^1) - \hat{\mu}_i$ and the input values are denoted as $\underline{\mathbf{z}} = \left(\mathbf{z}, \mathbf{z}_1, \dots, \mathbf{z}_M\right)$. Different choices of the partitioning schemes for $\mathbf{z}_i$ correspond to estimators with different properties. It is possible to obtain estimators like MLMC [4] and MFMC [5], which are recursive estimators, and estimators like ACV-IS and ACV-MF [6], which are designed to remove the dependence on a hierarchical/recursive assumption.

For all the ACV in the estimator form, both the optimal weights $\underline{\alpha}^{\mathrm{ACV}}$ and the variance of $\tilde{Q}$ can be obtained as

$$\begin{aligned} \underline{\alpha}^{\mathrm{ACV}} &= -\mathbb{C}ov\left[\underline{\Delta}, \underline{\Delta}\right]^{-1} \mathbb{C}ov\left[\underline{\Delta}, \hat{Q}\right] \\ \mathbb{V}ar\left[\tilde{Q}\left(\underline{\alpha}^{ACV}\right)\right] &= \mathbb{V}ar\left[\hat{Q}\right]\left(1 - \mathbb{C}ov\left[\underline{\Delta}, \hat{Q}\right]^{\mathrm{T}} \frac{\mathbb{C}ov\left[\underline{\Delta}, \underline{\Delta}\right]^{-1}}{\mathbb{V}ar\left[\hat{Q}\right]} \mathbb{C}ov\left[\underline{\Delta}, \hat{Q}\right]\right) \\ &= \mathbb{V}ar\left[\hat{Q}\right]\left(1 - R^2_{\mathrm{ACV}}\right). \end{aligned} \tag{6}$$

Partitioning schemes that are used to obtain several estimators are reported in Table 1.

As demonstrated in [6], the advantage of estimators that do not rely on a recursive assumption is that they are able to obtain a variance reduction, in the limit of $r_i \to \infty$, that converges to R^2_{OCV}, which is the maximum theoretical variance reduction that can be achieved. On the other side, recursive estimators, even in the limit of $r_i \to \infty$, are limited by the the variance reduction of a control variate with only a single low-fidelity model, *i.e.* $1 - \rho_1^2$, which can severely limit the performance of a multi-model estimator.

Algorithm	Relation between $\mathbf{z}$ and $\mathbf{z}_i$	$\mathbf{z}_i^1$	$\mathbf{z}_i^2$
OCV	$\mathbf{z}_i = \mathbf{z}$	$\mathbf{z}$	$\emptyset$
MLMC [4]	$\mathbf{z}_1^1 = \mathbf{z}, \mathbf{z} \cap \mathbf{z}_i = \emptyset$ for $i > 1$	$\mathbf{z}_{i-1}^2$	$\mathbf{z}_i \setminus \mathbf{z}_i^1$
MFMC [5]	$\mathbf{z}_i \supset \mathbf{z}$ for all i	$\mathbf{z}_{i-1}$	$\mathbf{z}_i$
ACV-IS	$\mathbf{z} \cap \mathbf{z}_i = \mathbf{z}_i^1, \mathbf{z}_i^2 \cap \mathbf{z}_j^2 = \emptyset$ for $1 \leq i \neq j$	$\mathbf{z}$	$\mathbf{z}_i \setminus \mathbf{z}_i^1$
ACV-MF	$\mathbf{z} \cap (\mathbf{z}_i \setminus \mathbf{z}_i^1) = \emptyset$	$\mathbf{z}$	$\mathbf{z}_i$

Table 1 Sampling schemes for ACV estimators from [6].

III. Graph Selection

While the non-recursive ACV-MF and ACV-IS estimators are optimal in the limiting case of $r_i \to \infty$, it has been illustrated in [6] that the recursive estimators can lead to superior variance reduction in the small r_i regime. In these small r_i regimes, the recursive structure accelerates the convergence toward the variance reduction of a control variate with only a single low-fidelity model. Thus, it is clear that the variance reduction of an ACV estimator is a function of the structure of that ACV estimator, and that the optimal structure may be dependent upon available data, among other factors.

With the aim of investigating optimal ACV structure, generalized versions of ACV-MF and ACV-IS were introduced in [7], dubbed GMF and GIS, respectively. These generalized ACV estimators allow for the restructuring of the estimators based on a graph that describes the control variate structure. The sampling schemes for GMF and GIS are detailed in Table 2. The control variate graphs are defined by a series of directed connections $(i) \to (j)$ that indicate that Q_i acts as a control variate of Q_j *. All zero-rooted, acyclic graphs in which each node has a single directed connection represent valid ACV structures. The control variate graphs for the recursive estimators (i.e. MLMC and MFMC) are represented by an ordered series of connections, whereas the graph for ACV-IS and ACV-MF is a flat structure (see Figure 1 for an example).

Algorithm	$\mathbf{z}_i^2$	Definition of $(i) \to (j)$
GMF	$\mathbf{z}_i^2 \subset \mathbf{z}_j^2$ or $\mathbf{z}_i^2 \supset \mathbf{z}_j^2$ for all i, j	$\mathbf{z}_i^1 = \mathbf{z}_j^2$
GIS	$\mathbf{z}_i^2 = \mathbf{z}_i^1 \cup \mathbf{z}_i^{2'}$ for all i, $\mathbf{z}_i^{2'} \cap \mathbf{z}_j^{2'} = \emptyset$ for $i \neq j$	$\mathbf{z}_i^1 = \mathbf{z}_j^{2'}$

Table 2 Sampling schemes for generalized ACV estimators from [7].

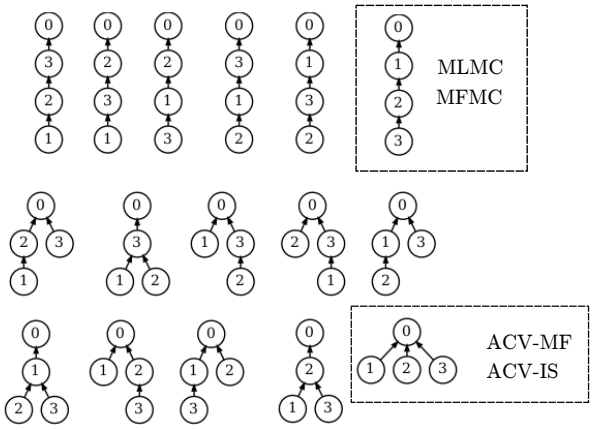

Fig. 1 All valid control variate graphs for $M = 3$.

*In the context of the model structure graphs, the high-fidelity model Q is referenced with the 0 index i.e. $Q_0 = Q$. Similarly, $\mathbf{z}_0 = \mathbf{z}_0^1 = \mathbf{z}_0^2 = \mathbf{z}$.

It was illustrated in [7], that for a reasonable number of low-fidelity models, the optimal control variate graph can be found through enumeration. In the most robust implementation of this enumeration, all subsets of low-fidelity models are also considered. However, in the case of large numbers of possible low-fidelity models, the large number of control variate graphs can make enumeration intractable. Thus, the enumeration approach is not suited for the case of low-fidelity model definitions by continuously-variable hyperparameters. Ideally, these types of model selection and tuning are made independent of graph selection. In this work we will focus on the hyperparameters tuning for a fixed graph and we will also show preliminary results on the graph selection, given different choices of the hyperparameters' values.

IV. Model Selection and Tuning

For all the algorithms written in an ACV form, the total cost C^{tot} is the sum of the evaluations for the low- and high-fidelity models:

$$C^{tot} = N\left(w + \sum_{i=1}^{M} w_i r_i\right), \tag{7}$$

where w and w_i represent the computational burden that is necessary to obtain an evaluation of Q and Q_i, respectively. As customary in multi-model UQ, the estimator variance can be minimized subject to the condition that the total cost, expressed as equivalent number of high-fidelity simulations, is less than or equal to a prescribed target, *i.e.*

$$C^{tot,eq} \leq \frac{C^{tot}}{w} = N\left(1 + \sum_{i=1}^{M} \frac{w_i}{w} r_i\right). \tag{8}$$

In this work, we are interested in augmenting this optimization problem by considering an additional aspect, *i.e.* the selection of the hyperparameters in each of the low-fidelity models such that the variance reduction is maximized, for a given cost. Examples of the importance of the hyperparameters tuning have been recently provided in [11, 12] for applications relative to simple CFD inspired problems with Reduced Order Models (ROMs) and computer networks, respectively. However, in the present work, we integrate the model tuning in the optimization and we derive a numerical solution that allows for the tuning and the sample allocation, at the same time. For simplicity of exposure, we focus on ACV-MF, although the approach can be adopted for any ACV-like estimator introduced in the previous section. The formulation of the optimization problem we want to solve takes the following form:

$$\underset{N,\underline{r},\underline{\beta}}{\arg\min} \frac{\mathbb{V}ar[Q]}{N}\left(1 - R^2_{ACV-MF}(\underline{r},\underline{\beta})\right) \quad \text{s.t.} \quad C^{tot,eq}\left(N,\underline{r},\underline{\beta}\right) \leq C_{target}, \tag{9}$$

where $\underline{r} = [r_1, \ldots, r_M]^{\mathrm{T}}$ is the vector of oversampling ratios for each of the M low-fidelity models, *i.e.* $r_i = \mathrm{card}(\underline{\mathbf{z}}_i)/N$, and $\underline{\beta} = \left[\underline{\beta}_1, \ldots, \underline{\beta}_M\right]^{\mathrm{T}}$ represents the set of hyperparameters for the low-fidelity models. We note here that each low-fidelity model can, in principle, be controlled by more than one hyperparameter, and therefore it is represented by a vector of hyperparameters $\underline{\beta}_i$ whose cardinality is $h_i \geq 0$, where $h_i = 0$ corresponds to the case of a pre-assigned, *i.e.* non-tunable, low-fidelity model.

A few considerations are in order here. First, the variance reduction term R^2_{ACV-MF} is only impacted by the selection of $\underline{r}$ and $\underline{\beta}$, since it can be evaluated by using the following expression

$$R^2_{ACV-MF} = \mathbf{a}^{\mathrm{T}}\left[\mathbf{C} \circ \mathbf{F}\right]^{-1}\mathbf{a} = \frac{1}{\mathbb{V}ar[Q]}\left[\mathrm{diag}\left[\mathbf{F}(\underline{r})\right] \circ \mathbf{c}(\underline{\beta})\right]^{\mathrm{T}}\left[\mathbf{C}(\underline{\beta}) \circ \mathbf{F}(\underline{r})\right]^{-1}\left[\mathrm{diag}\left[\mathbf{F}(\underline{r})\right] \circ \mathbf{c}(\underline{\beta})\right], \tag{10}$$

where we indicated with $\circ$ the Hadamard, or element-by-element, product. We also note that the matrix $\mathbf{F}$ encodes the particular sampling strategy, *e.g.* ACV-MF in this case, and it only depends on $\underline{r}$, whereas the entries of $\mathbf{c}$ and $\mathbf{C}$ depend on the covariances among all the models, and, therefore, they depend solely on $\underline{\beta}$. In a practical application, it is often, if not always, impossible to determine *a priori* and in closed-form the relationships $\mathbf{c} = \mathbf{c}(\underline{\beta})$ and $\mathbf{C} = \mathbf{C}(\underline{\beta})$. This challenge can be, at least notionally, overcome by resorting to a surrogate for each entry of $\mathbf{c}$ and $\mathbf{C}$. The particular details of this surrogate construction are application specific, and we will provide more details in the next sections. Moreover, we are interested in considering the general case of a model cost, for each low-fidelity model, being impacted by the vector of hyperparameters, *i.e.* $w_i = w_i(\underline{\beta}_i)$. Differently from the entries of $\mathbf{c}$ and $\mathbf{C}$, the relationship between the

cost of a model and the hyperparameters is not always unknown; see for instance [11] in which, by using linear algebra manipulations, the authors determined the cost of a ROM model as a function of the number of basis functions and time step. In general, we can argue that for models defined with hyperparamters that affect the spatial and/or temporal resolutions, it is possible to estimate the cost relationships beforehand. In all other cases, the cost can be treated as an additional quantity to be recovered via a surrogate.

Finally, we note that, once the optimal solution for N, $\underline{r}$ and $\underline{\beta}$ is obtained, the optimal weights can be evaluated by resorting to Eq. (6), hence the ACV weights $\underline{\alpha}$ do not appear explicitly in the optimization problem. In the next section, we show how these concepts can be used to improve the performance of an ACV estimator for a simple analytical problem.

A. An analytical demonstration problem

Before deploying the model tuning approach to the trajectory re-entry problem, we will provide an analytical demonstration that illustrates the features and potential of the proposed approach. We consider the model problem discussed in [6], in which two low-fidelity models are used to improve the precision of the mean estimator. The high- and low-fidelity models can be defined as

$$\begin{aligned} Q &= \sqrt{11}y^5 \\ Q_1 &= \sqrt{7}\left(\cos\theta_1 x^3 + \sin\theta_1 y^3\right) \\ Q_2 &= \sqrt{3}\left(\frac{\sqrt{3}}{2}x + \frac{1}{2}y\right), \end{aligned} \tag{11}$$

where $x, y \sim \mathcal{U}(-1, 1)$. These models have a unitary variance and therefore the correlation and covariance matrices are equivalent and defined in Table 3.

In this formulation, the properties of this set of models change according to the selection of the parameter θ_1 only, which we assume here can be optimized, such that $\underline{\beta} = \theta_1$. For instance, in a realistic application θ_1 could represent the number of basis terms in a ROM or a continuous parameter in a physical model, *e.g.* a coefficient for a turbulent model in a Reynolds-Averaged Navier-Stokes model. In the case of EDL trajectory simulation, this could be a time integration step in a reduced integration model or the hyperparameters of a machine learning surrogate. For illustrative purposes, in this example, we assume a cost law that, over a design range $\pi/6 \leq \theta_1 \leq \pi/2$, produces a cost, for the first low-fidelity model Q_1, that varies between the cost of evaluating the model Q_2 (the coarsest model) and Q (the high-fidelity) according to the law

$$\log w_1 = \log w_2 + \frac{\log w_2 - \log w}{\theta_2 - \theta}\left(\theta_1 - \theta_2\right) \tag{12}$$

using $\theta = \pi/2$ and $\theta_2 = \pi/6$.

This law for the cost mimics the case of a hyperparameter that can make the Q_1 more correlated with the high-fidelity model, at the price of a sharp increase of the model cost. This assumption is justified by the fact that as θ_1 approaches θ, the model is increasingly similar to the high-fidelity and, consequently, its cost rapidly becomes similar to the high-fidelity. For this example, we assume $w = 1$ and $w_2 = 10^{-3}$, which correspond to a lowest fidelity model three orders of magnitude cheaper than the high-fidelity, to evaluate. Similarly, by construction, we can expect the model Q_1 to correlate better with Q for $\theta_1 \to \theta$. The cost w_1 and correlation properties of the Q_1 model, with both the high-fidelity ρ_1 and lowest fidelity ρ_{12}, are reported in Figure 2 for the entire range of θ_1 values considered.

As evident from the results reported in [6], whenever the cost w_1 is not affected by θ_1, the optimal estimator is obtained in the correspondence of the θ_1 that maximizes the correlation ρ_1, *i.e.* $\theta_1 = \theta$. However, in our current scenario

	Q	Q_1	Q_2
Q	1	$\frac{\sqrt{77}}{9}\sin\theta_1$	$\frac{\sqrt{33}}{14}$
Q_1	*sym*	1	$\frac{\sqrt{21}}{10}\left(\sin\theta_1 + \sqrt{3}\cos\theta_1\right)$
Q_2	*sym*	*sym*	1

Table 3 Correlation/covariance matrix for analytical test problem.

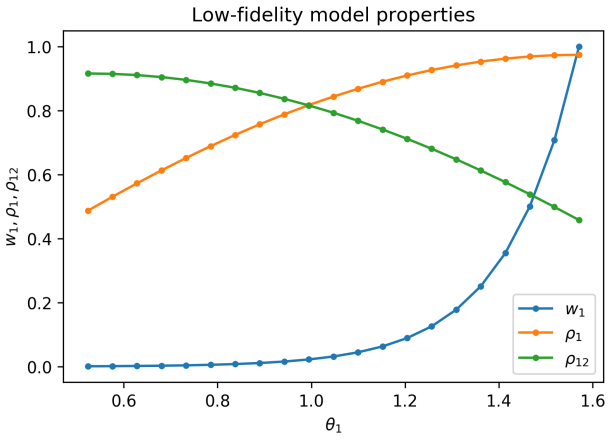

Fig. 2 Low-fidelity properties for the test case with hyperparameter θ_1.

the situation is more complex and a non-trivial trade-off between correlation ρ_1, ρ_{12} and the cost w_1 is present. This optimization problem needs to be solved numerically, since a closed form solution is not going to be available in general. In Figure 3, to illustrate this trade-off, for several values of θ_1, the variance for the ACV-MF estimators is reported. We note here that this is a simplified exploratory study which does not require solving the coupled optimization problem. Instead, the sample allocation and resulting estimator variance for ACV-MF is solved multiple times for a predetermined set of values for θ_1 and the optimal value is selected from the results. The optimization problem is solved for a total cost $C_{target} = 1000$, which corresponds to a MC variance equal to $\mathbb{V}ar\left[\hat{Q}\right] = 10^{-3}$ (since $\mathbb{V}ar[Q] = 1$ for construction).

This simple example serves the purpose of illustrating how the optimization of the hyperparameters might be needed to obtain an efficient multi-model estimator, since choices that are far from the optimal region, despite being still effective with respect a single fidelity MC estimator, could only achieve a marginal efficiency gain.

In the next section we will demonstrate this approach on the trajectory re-entry problem and we will illustrate the impact of the pilot sample statistics in the overall error and reliability of this coupled optimization formulation.

V. ADEPT Case Study

Adaptable Deployable Entry and Placement Technology (ADEPT) is a deployable aeroshell which can be used to reduce thermal loads and protect payloads during entry into a planetary atmosphere. In September 2018, a test flight of ADEPT was performed to demonstrate ADEPT's suitability for future missions [13]. In previous studies of ADEPT SR-1, pre-flight trajectory simulations were compared to the SR-1 flight data using standard MC and multi-model MC [2, 3]. Some of the multi-model MC analysis from the work of Warner et al. [3] is revisited in this section through the lens of model tuning.

Program to Optimize Simulated Trajectories II (POST2), a legacy code developed at NASA Langley Research Center, is used here for trajectory simulation. Notably, POST2 is highly configurable and allows for modification of many hyperparameters which affect the fidelity of trajectory simulations. This work focuses on two such hyperparameters: the integration time step Δt and the atmosphere model used. The high-fidelity model Q uses $\Delta t = 10^{-3}$s and a *full-physics* atmosphere model: Earth Global Reference Atmospheric Model (GRAM) [14]. Two lower fidelity models are considered: a mid-fidelity model, Q_1, using the same high-fidelity atmosphere model but with reduced integration (i.e., increased Δt) and a low-fidelity model, Q_2, that uses a simplified atmosphe model, henceforth referred to as *reduced-physics*. In particular, Q_2 uses U.S. 1976 Standard Atmosphere conditions and neglects any calculations pertaining to winds, which results in roughly a $4X$ speedup. The time steps for each of these models (Δt_1 and Δt_2, respectively) may be varied, and thus represent the model tuning hyperparameters.

For all the subsequent analyses, we consider a dataset of 25 time steps equally spaced in logarithm space between 0.001 and 0.25 seconds. These time steps correspond to a uniform (in logarithm space) tessellation of the space of hyperparameters Δt_1 and Δt_2. For each hyperparameter location, we evaluated each model with a total of 500 random samples, shared among all hyperparameters locations. These data are used to extract the statistics needed for the different optimization scenarios, which will also serve as reference solution.

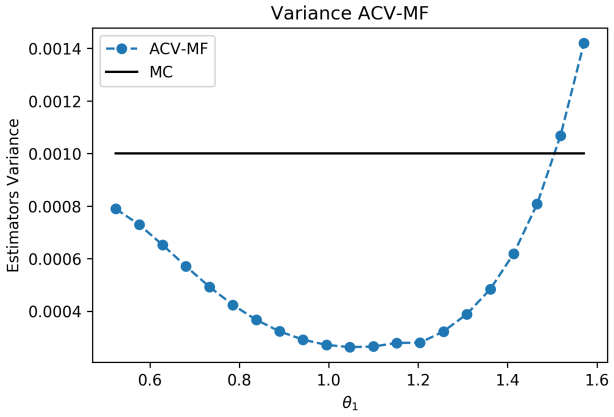

Fig. 3 Variance of the ACV-MF estimator for different choices of the hyperparameter θ_1. For an optimal value of the hyperparameter θ_1, the ACV-MF estimator achieves a variance that is approximately 5 times less than the MC estimator with an equivalent computational cost. However, for values of θ_1 that deviate from the optimal value, the efficiency of ACV-MF is drastically reduced and its variance can even be larger than the one of MC.

A. Example 1: ACV-MF tuning for the middle fidelity model

In a first example, a sample allocation optimization is performed considering model tuning for a single parameter. In this case the low-fidelity time-step is fixed to $\Delta t_2 = 0.25$s and the mid-fidelity time step is the sole tuning parameter, *i.e.* $\underline{\beta} = \Delta t_1 \in \Delta t_1^{tun}$. We consider the total time to touchdown as the QOI for this analysis. Two possible analysis scenarios are considered for the optimization: the middle model can be either the full or the reduced physics model. The model definitions are shown in Table 4. The approximate normalized costs w_i/w for evaluating the models, in this first scenario, are illustrated in Figure 4. The correlation coefficients for the models are reported, for reference, in Figure 5 for all the time steps considered in both scenarios, with either full or reduced physics.

Due to the regular convergence rate of the two model forms, either full or reduced, we assume to know in closed-form the relationship for the cost of the first LF model, *i.e.* $w_1 = w_1(\Delta t_1)$, see Figure 4. For this application, this can be easily constructed beforehand with few pilot batches in the tuning space. By using this assumption, the optimization formulation in Section IV can be further simplified as

$$\underset{N,r_1,r_2,\Delta t_1}{\arg\min} \; \frac{\mathbb{V}ar[Q]}{N}\left(1 - R^2_{ACV-MF}(r_1, r_2, \Delta t_1)\right) \quad \text{s.t.} \quad N\left(1 + r_1\frac{w_1(\Delta t_1)}{w} + r_2\frac{w_2}{w}\right) \leq C_{target}. \tag{13}$$

For this analysis we consider a total computational budget of 1000 equivalent HF simulations, *i.e.* $C_{target} = 1000$.

We now focus on the expression of R^2_{ACV-MF}, which is a function of the oversampling ratios (r_1 and r_2) and the hyperparameter Δt_1. For the ACV-MF, in this analysis scenario with two low-fidelity models we can simplify the expression introduced in Section IV as

$$R^2_{ACV-MF} = \frac{q_1(r_1)q_2(r_2)}{q_1(r_1)q_2(r_2) - q_m^2(r_1,r_2)\rho_{12}^2(\Delta t_1)}\left(q_1(r_1)\rho_1^2(\Delta t_1) + q_2(r_2)\rho_2^2 - 2q_m(r_1,r_2)\rho_1(\Delta t_1)\rho_2\rho_{12}(\Delta t_1)\right), \tag{14}$$

where $q_1 = \frac{r_1-1}{r_1}$, $q_2 = \frac{r_2-1}{r_2}$ and $q_m = \frac{\min(r_1,r_2)-1}{\min(r_1,r_2)}$ are used as shorthands for the elements of the matrix $\mathbf{F}$, while ρ_1, ρ_2 and ρ_{12} indicate the correlation between Q-Q_1, Q-Q_2 and Q_1-Q_2, respectively. From this expression it is possible to

	Example 1		Example 2	
	Physics	Time step (s)	Physics	Time step (s)
Q	Full	0.001	Full	0.001
Q_1	Full/Reduced	$0.001 \leq \Delta t_1 \leq 0.25$	Full	$0.001 \leq \Delta t_1 \leq 0.25$
Q_2	Reduced	0.25	Reduced	$0.001 \leq \Delta t_1 \leq 0.25$

Table 4 Model definitions for ADEPT case study. QOI: total time to touchdown

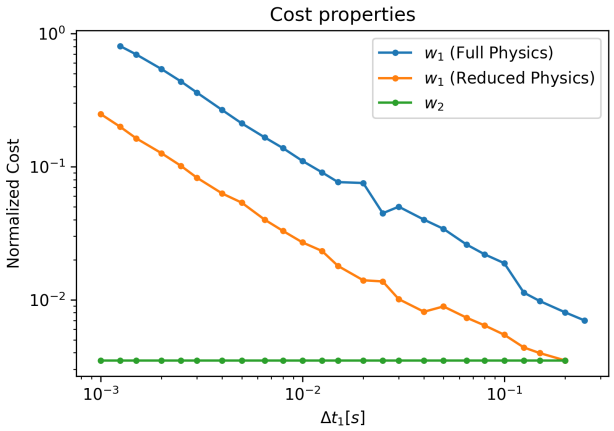

Fig. 4 ADEPT SR-1 model evaluation costs using POST2 trajectory simulator.

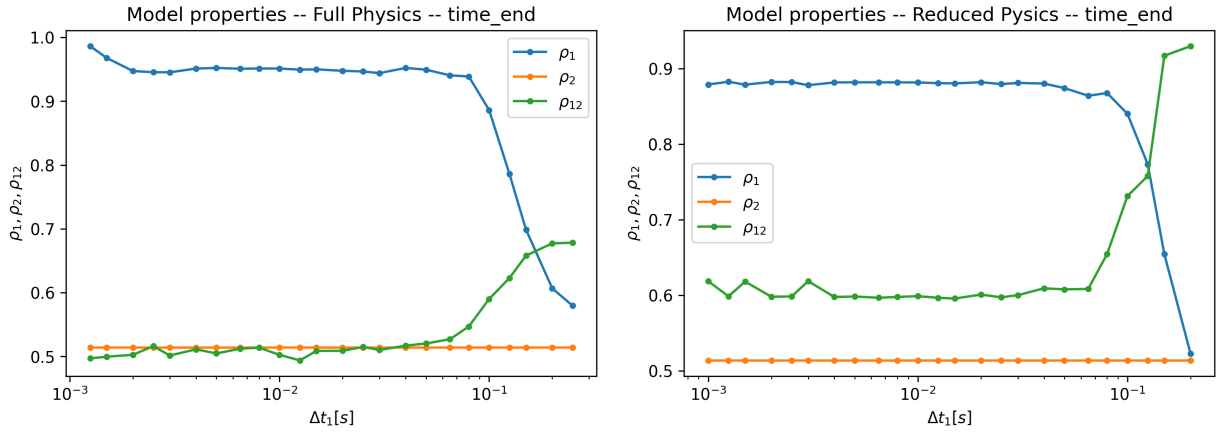

Fig. 5 Correlations of models for ADEPT Example 1 estimated using 500 samples at each time step for both the full and reduced physics case.

note that, in our analysis scenarios in which the coarsest model is fixed, we still need to estimate all the correlations that depend on the first LF model, *i.e.* ρ_1 and ρ_{12}. These quantities will need to be reconstructed via surrogates.

We are interested here in performing a comparison between the reference solutions we can obtain by using the entire dataset and an optimization problem run starting from a limited dataset, in both the hyperparameter space (N_{tun}) and number of samples (N_{pilot}). Once the batches of pilot runs are evaluated, surrogates are constructed for the correlations and used for the solution of the minimization problem, Eq. (13). For the surrogates, we consider a local quadratic polynomial and for the optimization we employ the `SLSQP` method provided in the python library `scipy`, and the following initial conditions, for the optimization, are provided: $\Delta t_1^0 = (\max \Delta t_1^{tun} + \min \Delta t_1^{tun})/2$, $r_i = \sqrt{\frac{w_i(\Delta t_1^0)}{w} \frac{\rho_i^2(\Delta t_1^0)}{1-\rho_i^2(\Delta t_1^0)}}$ and $N = \frac{C_{target}}{\left(1 + \sum_{i=1}^{M} \frac{w_i}{w} r_i\right)}$.

We consider a set of pilot samples corresponding to $N_{pilot} = [50, 100, 200]$ and $N_{tun} = [6, 12, 24]$ time steps. For each optimization study, we consider 200 repetitions in order to obtain statistics for the solutions. In Figure 6, we present the results achieved by the optimization for different combinations of number of samples and hyperparameters locations (N_{pilot}, N_{tun}), which are used to extract the correlations used for the construction of the surrogates; this case uses the full physics for the middle low-fidelity model. The results of the optimizations are reported as red crosses, while the green squares indicate the corresponding results that we would have achieved if the surrogates were known based on all the available samples in the dataset $(N_{pilot}, N_{tun}) = (500, 24)$. The reference performance of the ACV-MF and MC are also reported to illustrate the potential variance reduction achievable and the error introduced by all components of the optimization workflow.

A few consideration are in order here. For all the combinations of (N_{pilot}, N_{tun}) that we considered, reasonable performances can be achieved and the region of Δt_1^{tun} in which the efficiency of the ACV-MF algorithm is maximized, is correctly identified. The estimated performances of the algorithm, which are represented by the red crosses, are all close in magnitude, to the minimum variance achievable by ACV-MF in this scenario; however, as evident from observing the location of the green squares, the surrogates can introduce a significant error, which, in turn, can hamper the true variance reduction. This aspect will need to be considered more carefully when designing a full optimization algorithm based on iterative estimation of statistics and surrogates construction. Finally we note that in the vast majority of the cases, the ACV-MF algorithm can achieve an acceleration with respect to MC of about 2.5 to 5 times.

Since the construction of the surrogates is a key component of the approach, in Figure 7, we report the 200 realizations of the surrogates that are constructed for each tested combination of (N_{pilot}, N_{tun}). There are two noticeable trends. First, by moving from lower to higher values of pilot samples (left to right) we observe that, as expected, the statistical variability decreases. This is very evident when comparing the surrogates for the correlation ρ_{12}, which, for $N_{pilot} = 50$, results in values that overlap with the values of ρ_1. On the contrary, for $N_{pilot} = 200$, the overlap between ρ_{12} and ρ_1 is eliminated and the trends appear to more similarly follow the reference ones reported in Figure 5. By comparing the surrogates obtained starting from different numbers of hyperparameter locations, we observe that coarser datasets introduce low frequency oscillations, when compared with the surrogates constructed by using the full dataset of hyperparameter locations, *i.e.* $N_{tun} = 24$. In Figure 7, we also report with red dots, the final locations of the optimizations for both the surrogate predictions of ρ_1 and ρ_{12}. The dispersions of these dots around the reference surrogates values, represent the source of discrepancy between the estimated and reference variance reductions achieved by ACV-MF in Figure 6. We note here that in this demonstration we used a local surrogate, which allows for an interpolation of the available data. This approach is for demonstration purpose only since, in a study with a larger number of hyperparameters, a global surrogate should be preferred and overfitting should be avoided since the values are only estimated statistics from pilot and not their true values. This aspect again indicates that an iterative approach in which statistics and surrogates are iteratively refined in the course of the optimization should be adopted to control these courses of error.

We obtained qualitatively similar results for the optimization of the ACV-MF estimator with reduced physics for the middle model. These results are reported in Figure 8. We note here that the optimal region is shallower than the case with full physics for Q_1 and this is reflected in a larger variability of the optimal Δt_1 results. This is particularly evident in the results obtained for the coarsest dataset in N_{tun} (bottom row of Figure 8), in which three distinct values of Δt_1, which correspond to similar performance of ACV-MF, are obtained. For this case, we report the surrogates in Figure 9, where we note the same qualitative trends discussed for the previous case. However, a notable difference is that the local quadratic interpolation produces, in this case, values of ρ_{12} that artificially approach one. This effect is responsible for optimization results that are overestimated (in term of performance) when compared to the reference surrogates, as evident in the middle row of Figure 9, which corresponds to $N_{tun} = 12$. We note again that these challenges will

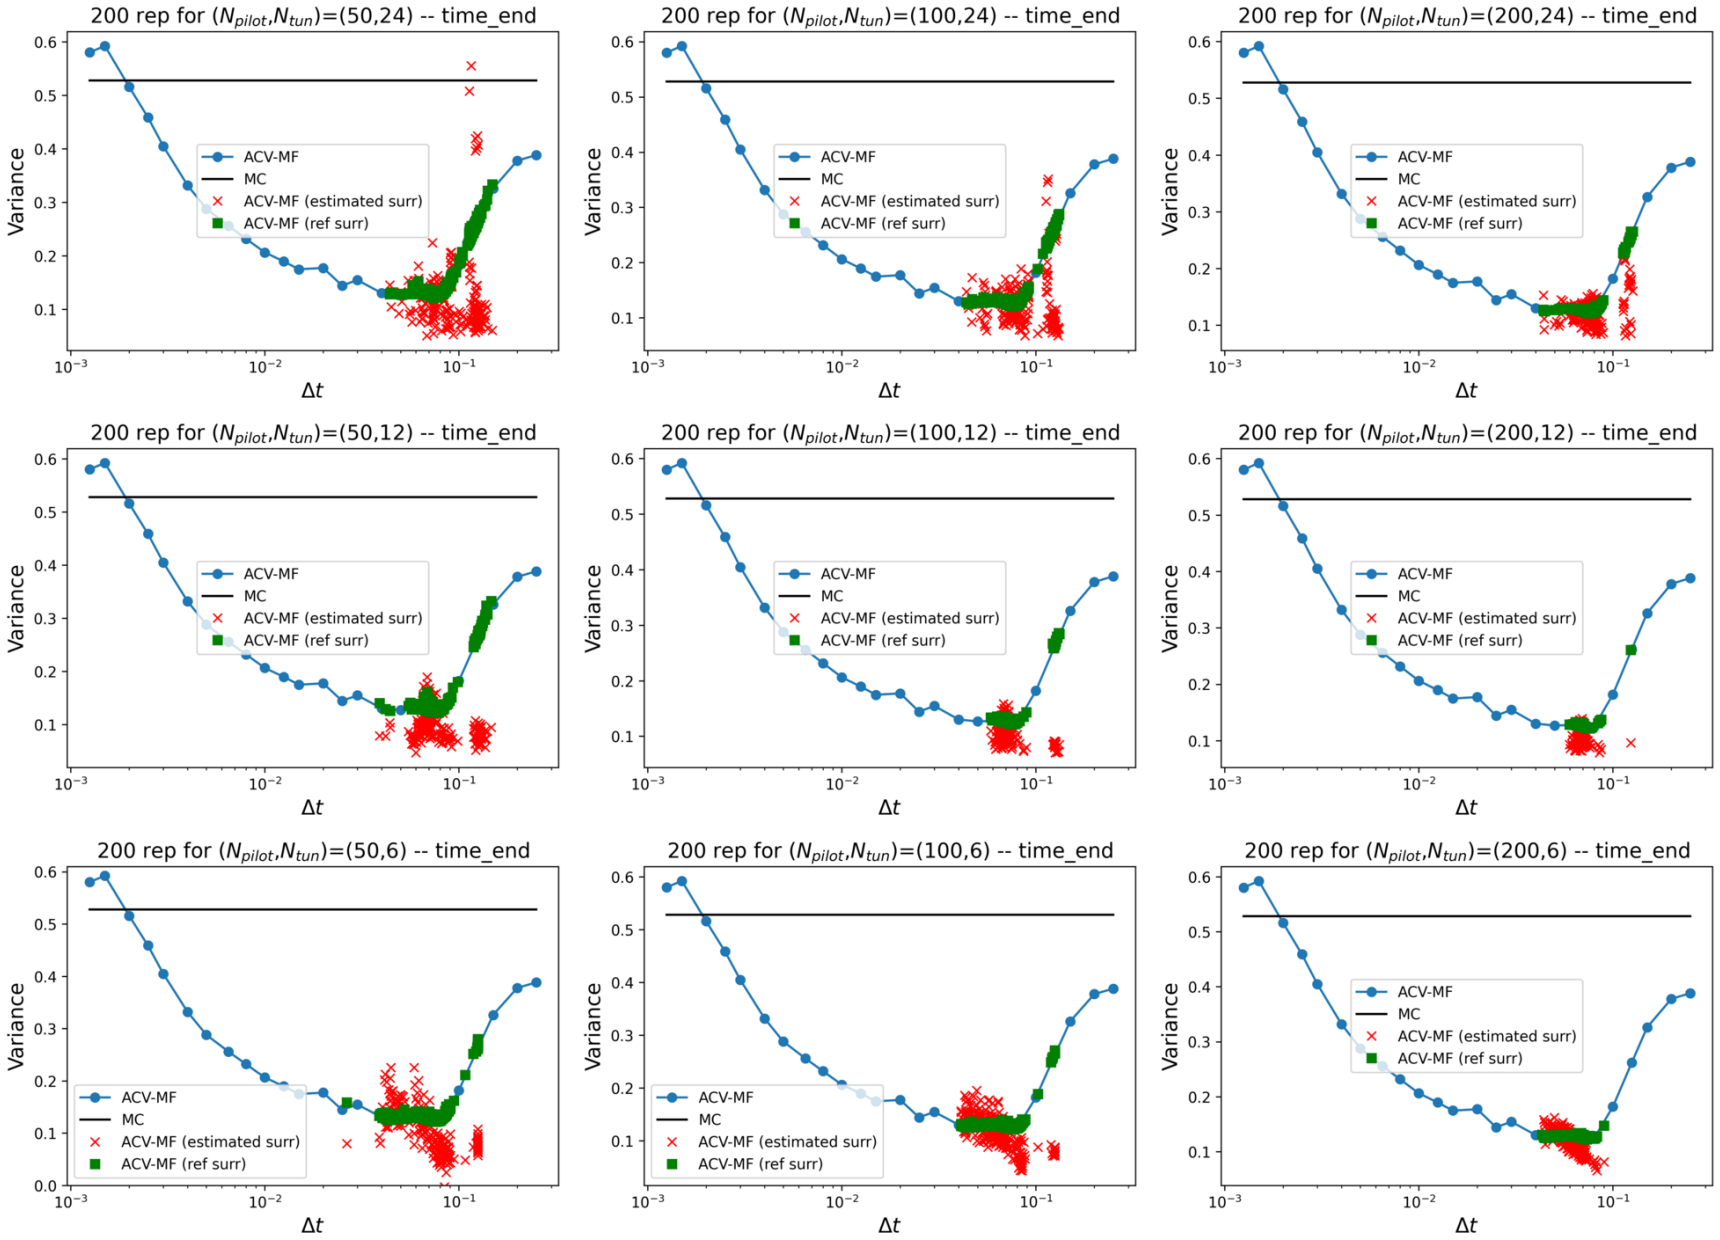

Fig. 6 Estimated variance for the QOI `time_end` with full order physics for the middle low-fidelity model of ACV-MF. The top, middle and bottom rows correspond to a decreasing number of hyperparameters locations for the pilot runs, $N_{tun} = 24, 12$ and 6, respectively. The left, middle and right columns correspond to an increasing number of samples for the pilot runs, $N_{pilot} = 50, 100$ and 200, respectively. The estimated variances of ACV-MF (red crosses) are all close, in magnitude, to the minimum variance achievable by ACV-MF in this scenario. However, from observing the ACV-MF variance corresponding to the reference surrogates (green squares), it is evident that the surrogates can introduce a significant error, which could hamper the achievable variance reduction.

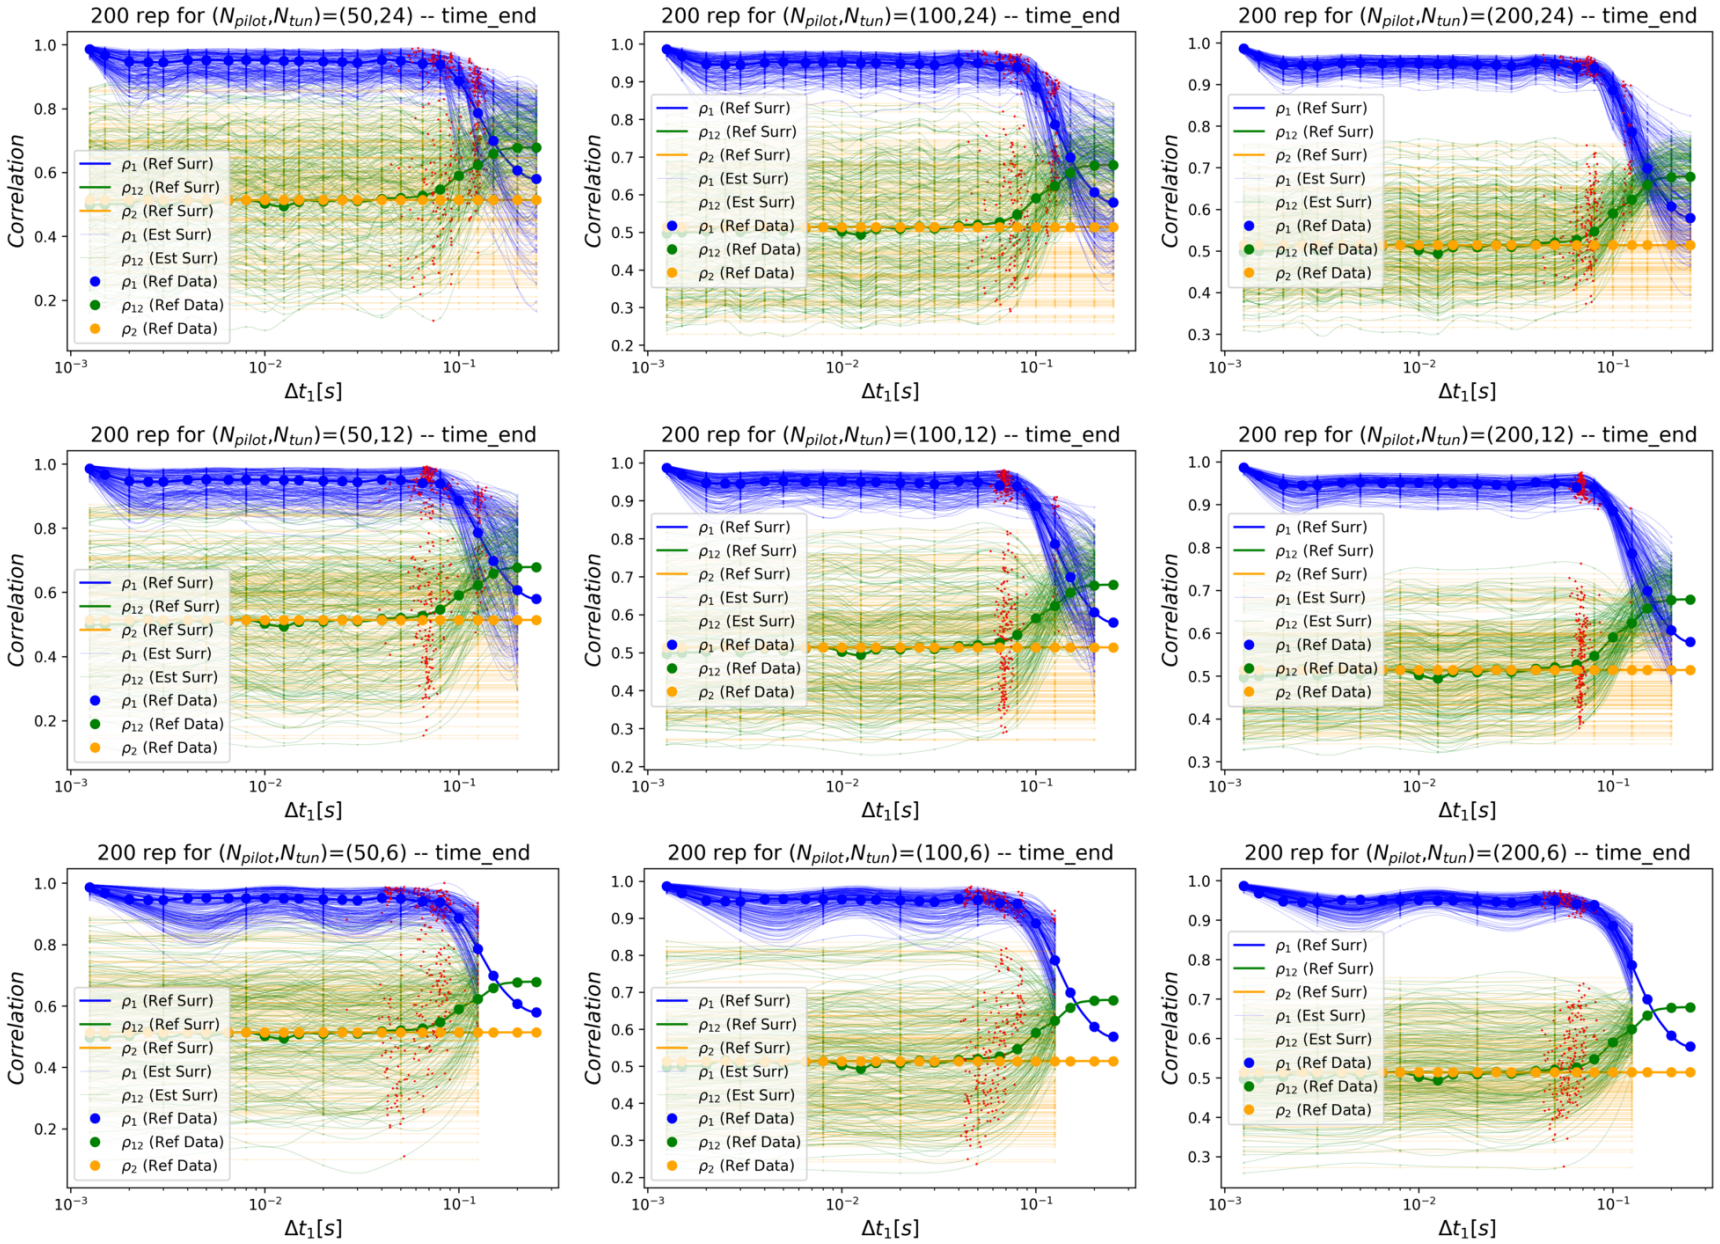

Fig. 7 Surrogates for the QOI `time_end` with full order physics for the middle low-fidelity model of ACV-MF. The top, middle and bottom rows correspond to a decreasing number of hyperparameters locations for the pilot runs, $N_{tun} = 24, 12$ and 6, respectively. The left, middle and right columns correspond to an increasing number of samples for the pilot runs, $N_{pilot} = 50, 100$ and 200, respectively. The red dots indicate the location in which the final optimization results terminate. Larger values of the pilot samples (N_{pilot}) reduce the statistical variability for both the surrogates and the optimization solution, while coarser hyperparameters datasets (N_{tun}) correspond to surrogates with lower spatial frequencies.

be mitigated by using global surrogates and their adaptive refinement over the course of the optimization whenever additional samples become available.

Finally, we point out that, despite the lower computational cost of the reduced physics model, the best performance are obtained by the ACV-MF estimator built upon the full physics as middle model, which can potentially achieve a higher variance reduction (approximately a factor 2).

B. Example 2: Exploratory performance study of ACV-MF performance with fully tunable LF models

A second example is considered here to (1) illustrate the potential of model tuning for multiple hyperparameters and (2) to illustrate the importance of graph selection. As in Example 1, the total touchdown time is used as the QOI. In this example however, the low-fidelity (reduced physics model) also has a tunable time step: $\underline{\beta} = \left\{\Delta t_1 \in \Delta t_1^{tun}, \Delta t_2 \in \Delta t_2^{tun}\right\}$ (see Table 4). Because there are two model tuning hyperparameters, the correlation functions become more complex. ρ_1 and ρ_2 are both one dimensional functions of Δt_1 and Δt_2, respectively. ρ_{12} is a two dimensional function of both Δt_1 and Δt_2. The higher dimension of these correlations relative to the first example corresponds to a higher burden on the surrogate training. In this example, we only provide the performance of the ACV-MF estimators as a function of all possible hyperparameter choices without performing the optimization, which is left for a subsequent study. The correlations functions, as function of the hyperparameters, are shown in Figure 10.

Using knowledge of the correlations, the variance of the ACV-MF estimator can be plotted as a function of the model hyperparameters Δt_1 and Δt_2 (see Figure 11). The variance reduction can change significantly depending on the ACV model graph that is chosen. The variance of the MFMC estimator (which differs from ACV-MF only in model graph) is illustrated in Figure 12a. Notably, the shape of the basin around the minimum variance is broader but with a sharper increase in variance as the correlation of the mid-fidelity drops (near $\Delta t_1 = 0.08$s). Figure 12b combines the variances of the ACV-MF and MFMC estimators into a single plot which illustrates the minimum variance (between the two estimators). The regions corresponding to each of the two estimators are divided by a white line. From this it can be seen that the optimal model graph can be a function of model tuning hyperparameters. In some cases it may be possible to enumerate over the possible model graphs and simply use the best estimator; however, as the number of models and/or number of model hyperparameters increases, enumeration becomes increasingly intractable. Thus an all-at-once optimization, which considers graph selection along with model tuning and sample allocation, may be beneficial. This aspect is out of the scope of the present work, but it will be addressed in a follow-up investigation.

VI. Conclusions and Perspectives

In this contribution, we proposed and investigated an integrated optimization approach that allows for the tuning and selection of low-fidelity models in the context of multi-model UQ. We have deployed and demonstrated the viability of a simplified instantiation of the proposed approach on a complex entry, descent, and landing application. Here, low-fidelity trajectory simulation models were tuned by selecting the integration time step that resulted in the minimum estimator variance. It was demonstrated that substantial reduction in variance for predicting time to touchdown for the ADEPT test flight could be achieved in this manner through optimization.

While the findings here indicate the potential advantages of model selection and tuning, the results are preliminary and also serve to identify challenges that could emerge in the design of a fully iterative procedure. Indeed, the proposed study is based on a single iteration from pilot batches to optimal hyperparameter identifications. Ideally, a complete algorithm that refines the underlying statistical surrogates should be based on multiple iterations. These improvements will be pursued in future work along with a production implementation of the approach in the Dakota UQ software.

The Dakota software tool kit has recently implemented numerical optimization-based solutions for multi fidelity Monte Carlo and approximate control variable estimation, in which the sample profiles can be iterated to achieve either accuracy or cost objectives. The suite of optimization solver options is being expanded from traditional gradient-based local optimization, e.g. Sequential quadratic programming (SQP) and Non-linear Interior Point (NIP) to include global and surrogate-based optimization. In the context of model tuning, surrogate-based optimization will be used to emulate changes in correlation and cost over the range of the tuning hyperparameters that are embedded within the set of low fidelity approximations. Here, we will rely on well-established capabilities for trust region model management that balance the efficient emulation of sparse and potentially noisy data with reliable convergence to the optimum of the datasets underlying these surrogates. In the case of mixed continuous and discrete hyper parameters, we plan to combine surrogate emulation with specialized Mixed Integer Nonlinear Programming (MINLP) approaches for solving the approximate sub-problems.

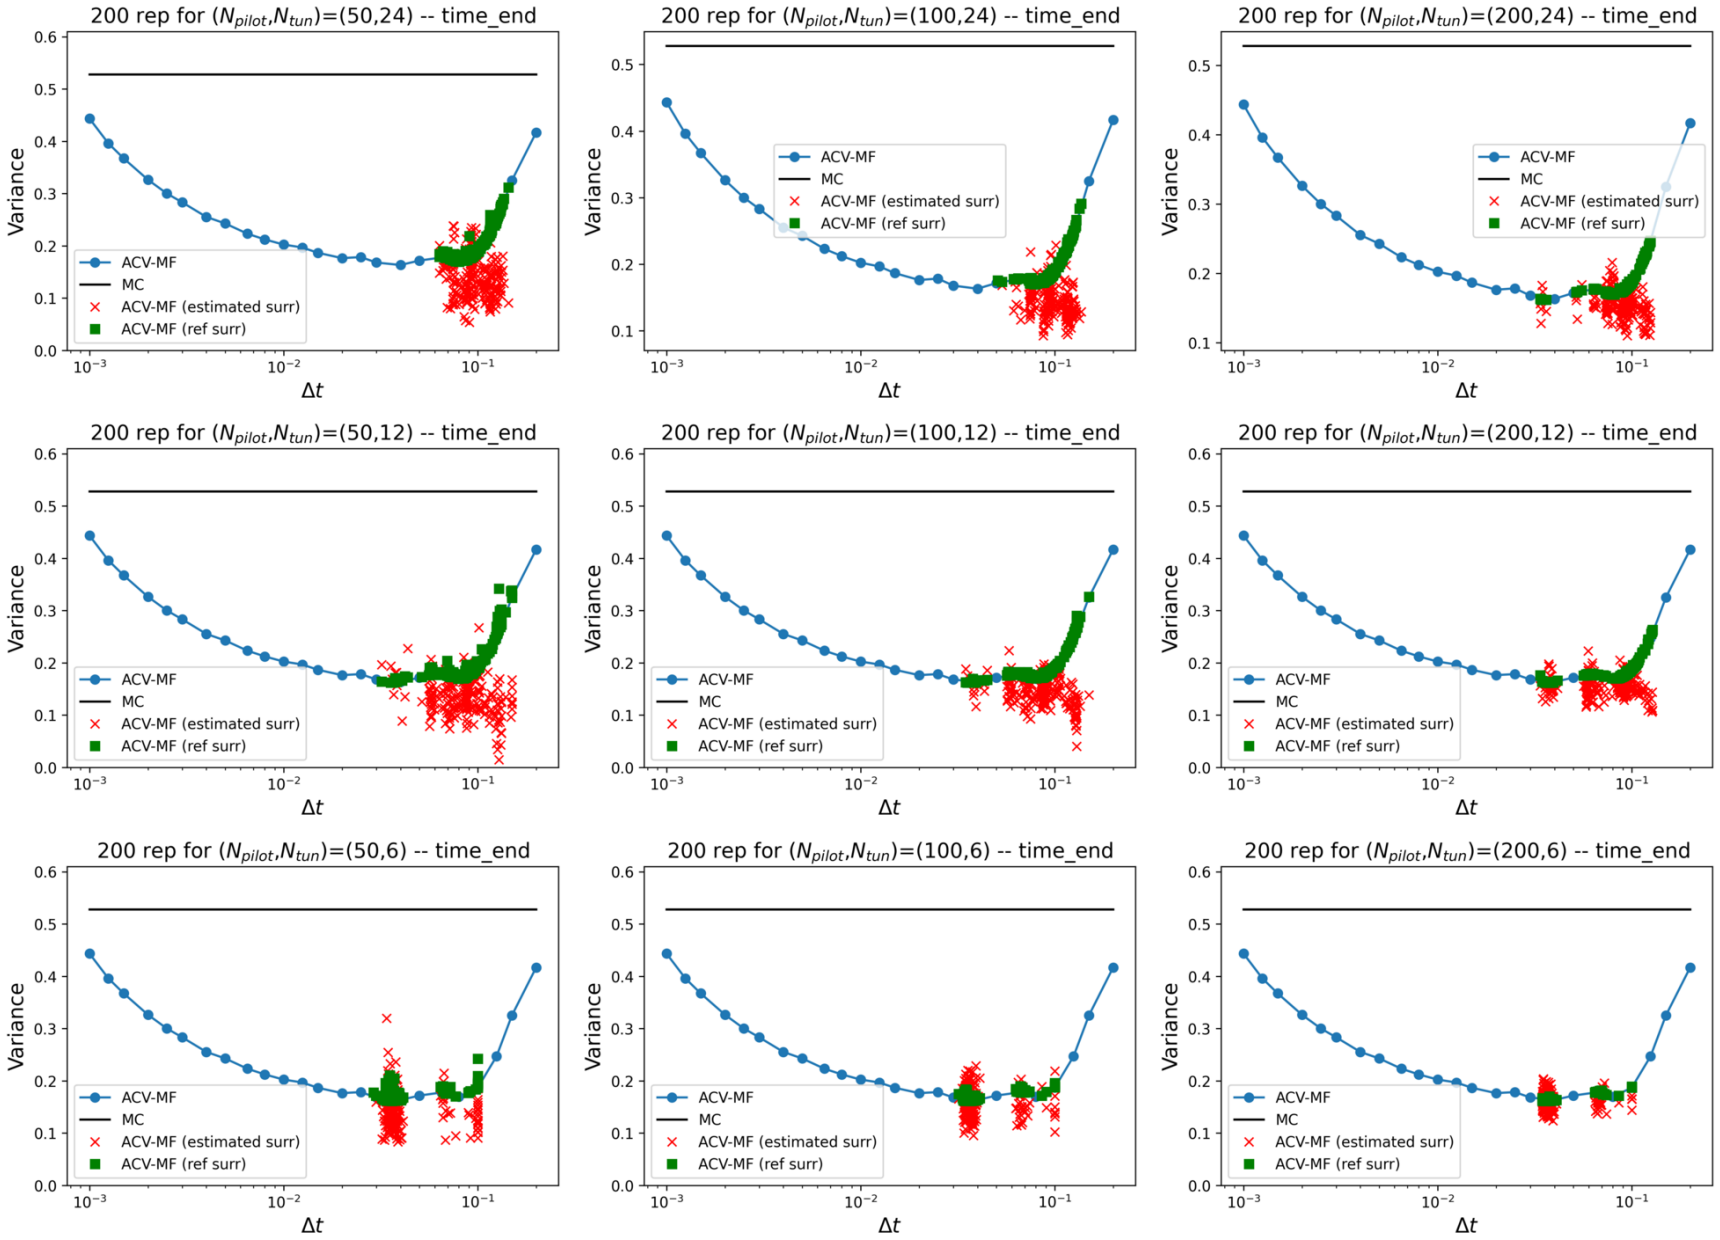

Fig. 8 Estimated variance for the QOI `time_end` with full order physics for the middle low-fidelity model of ACV-MF. The top, middle and bottom rows correspond to a decreasing number of hyperparameters locations for the pilot runs, $N_{tun} = 24, 12$ and 6, respectively. The left, middle and right columns correspond to an increasing number of samples for the pilot runs, $N_{pilot} = 50, 100$ and 200, respectively. The estimated variances of ACV-MF (red crosses) are all close, in magnitude, to the minimum variance achievable by ACV-MF in this scenario, however due to the shallow optimal region, there is an higher statistical variability for the optimization (Δt_1) results. For the coarsest dataset in N_{tun}, three distinct local values for Δt_1 correspond to similar values of the ACV-MF variance.

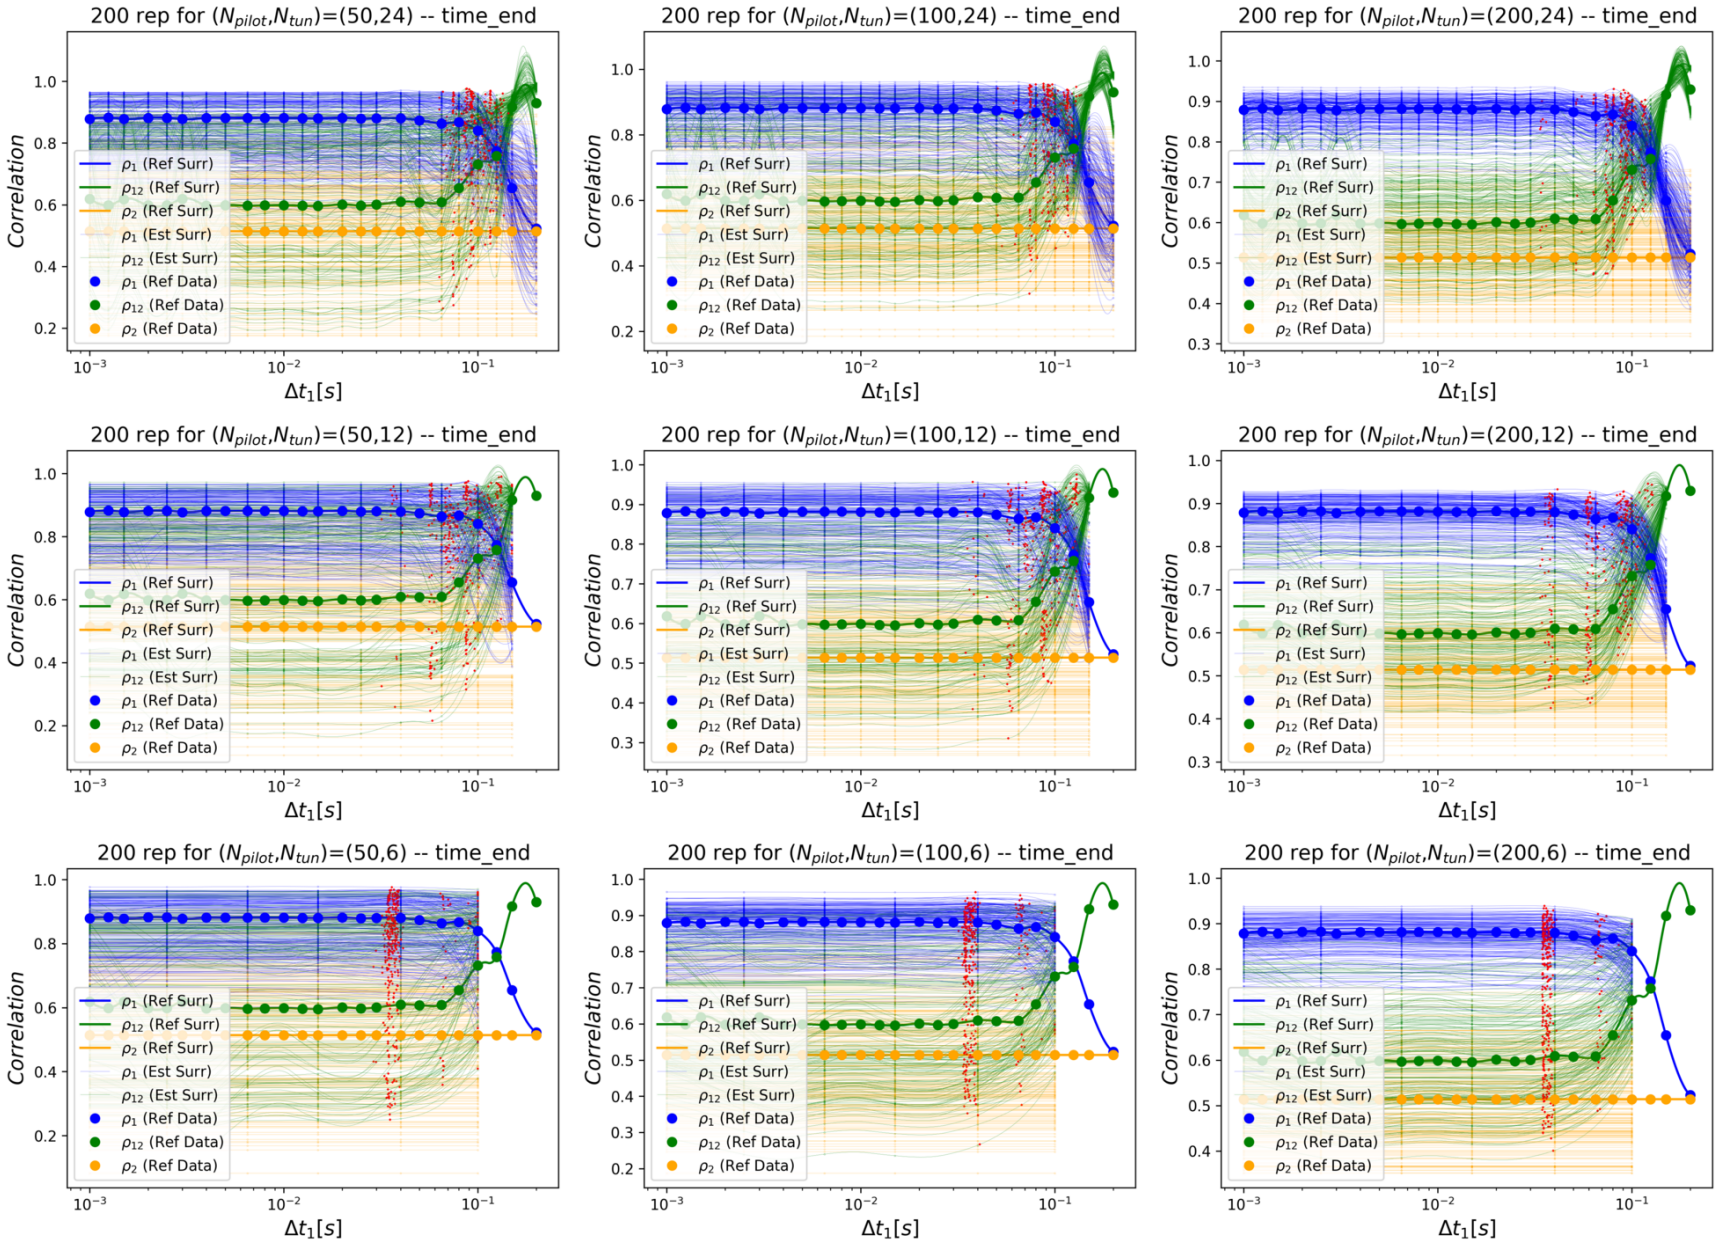

Fig. 9 Surrogates for the QOI `time_end` with full order physics for the middle low-fidelity model of ACV-MF. The top, middle and bottom rows correspond to a decreasing number of hyperparameters locations for the pilot runs, $N_{tun} = 24, 12$ and 6, respectively. The left, middle and right columns correspond to an increasing number of samples for the pilot runs, $N_{pilot} = 50, 100$ and 200, respectively. The red dots indicate the location in which the final optimization results terminate. We note that the local quadratic interpolation can produce values that artificially approach one, which correspond to an overestimation of the performance of the ACV-MF estimator, when compared to the reference surrogates values.

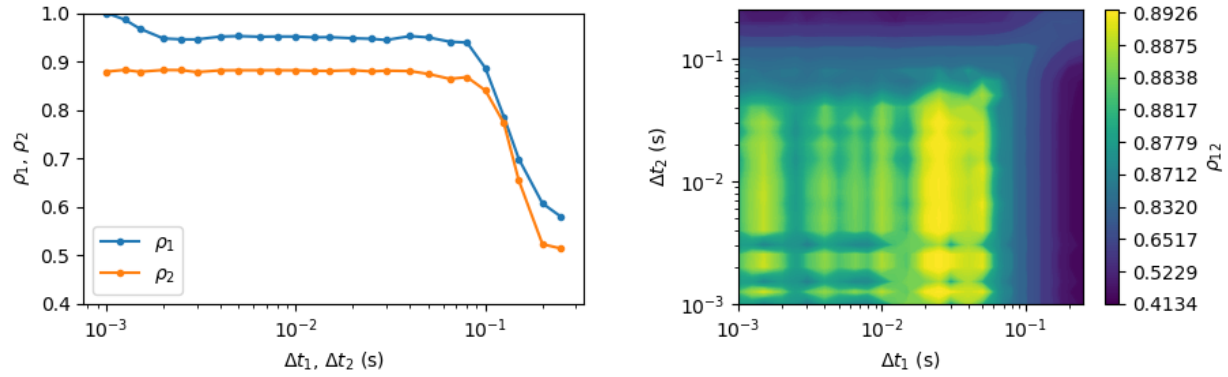

(a) **(b)**

Fig. 10 Correlations of models for ADEPT Example 2 estimated using 500 samples at each set of time steps.

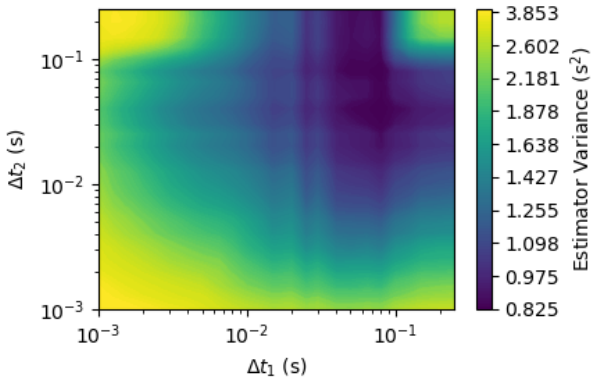

Fig. 11 Variance of ACV-MF estimator in ADEPT Example 2 for different choices of the hyperparameters Δt_1 and Δt_2. A significant variance reduction (approximately a factor 5) can be obtained for values of (approximately) $\Delta t_1 = 0.08$s and $\Delta t_2 = 0.03$s.

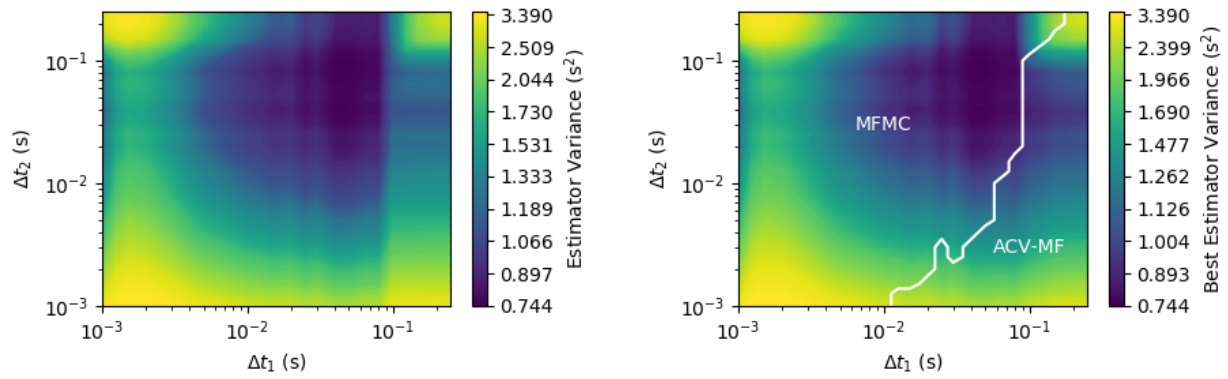

(a) **(b)**

Fig. 12 Variance of estimators in ADEPT Example 2 for different choices of the hyperparameters Δt_1 and Δt_2: (a) using the MFMC estimator and (b) minimum of ACV-MF and MFMC estimators. For this scenario, the most efficient estimator is obtained with MFMC in a region of the hyperparameters corresponding to approximately $\Delta t_1 = 0.06$s and $\Delta t_2 = 0.04$s. The best MFMC estimator has a variance which is approximately 90% of the ACV-MF variance.

Acknowledgments

Gianluca Geraci, John Jakeman, Michael Eldred and Alex Gorodetsky were supported by the Advanced Simulation and Computing (ASC) program and the Laboratory Directed Research Development (LDRD) program at Sandia National Laboratories. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. The views expressed in the article do not necessarily represent the views of the U.S. Department Of Energy or the United States Government.

References

[1] Cianciolo, A. D., and Powell, R. W., "Entry, Descent, and Landing Guidance and Control Approaches to Satisfy Mars Human Mission Landing Criteria," *27th AAS/AIAA Space Flight Mechanics Meeting*, AIAA, San Antonio, TX, 2017.

[2] Dutta, S., and Green, J. S., "Flight Mechanics Modeling and Post-Flight Analysis of ADEPT SR-1," *AIAA Aviation 2019 Forum*, AIAA, Dallas, TX, 2019. doi:10.2514/6.2019-2900.

[3] Warner, J., Niemoeller, S. C., Morrill, L., Bomarito, G., Leser, P., Leser, W., Williams, R. A., and Dutta, S., "Multi-Model Monte Carlo Estimators for Trajectory Simulation," *AIAA Scitech 2021 Forum*, 2021, p. 0761.

[4] Giles, M. B., "Multi-level Monte Carlo path simulation," *Operations Research*, Vol. 56, No. 3, 2008, pp. 607–617.

[5] Peherstorfer, B., Willcox, K., and Gunzburger, M., "Optimal Model Management for Multifidelity Monte Carlo Estimation," *SIAM Journal on Scientific Computing*, Vol. 38, 2016, pp. A3163–A3194. doi:10.1137/15M1046472.

[6] Gorodetsky, A., Geraci, G., Eldred, M., and Jakeman, J. D., "A generalized approximate control variate framework for multifidelity uncertainty quantification," *Journal of Computational Physics*, 2020, p. 109257. doi:https://doi.org/10.1016/j.jcp.2020.109257, URL `http://www.sciencedirect.com/science/article/pii/S0021999120300310`.

[7] Bomarito, G. F., Leser, P. E., Warner, J. E., and Leser, W. P., "On the Optimization of Approximate Control Variates with Parametrically Defined Estimators," *arXiv preprint arXiv:2012.02750*, 2020.

[8] Adams, B., et al., "Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.14 Reference Manual," Tech. rep., Sandia National Laboratories, Albuquerque, NM, Updated May 2021.

[9] Dalbey, K. R., Eldred, M. S., Geraci, G., Jakeman, J. D., Maupin, K. A., Monschke, J. A., Seidl, D. T., Swiler, L. P., Tran, A., Menhorn, F., , and Zeng, X., "Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.14 Theory Manual," Tech. Rep. SAND2020-12496, Sandia National Laboratories, Albuquerque, NM, Updated May 2021.

[10] Lavenberg, S. S., Moeller, T. L., and Welch, P. D., "Statistical results on control variables with application to queueing network simulation," *Operations Research*, Vol. 30, No. 1, 1982, pp. 182–202.

[11] Blonigan, P. J., Geraci, G., Rizzi, F., and Eldred, M. S., *Towards an integrated and efficient framework for leveraging reduced order models for multifidelity uncertainty quantification*, 2020. doi:10.2514/6.2020-0420, URL https://arc.aiaa.org/doi/abs/10.2514/6.2020-0420.

[12] Geraci, G., Crussell, J., Swiler, L., and Debusschere, B., "Exploration of multifidelity UQ sampling strategies for computer network applications," *International Journal for Uncertainty Quantification*, Vol. 11, No. 1, 2021, pp. 93–118.

[13] Tynis, J. A., Karlgaard, C. D., Williams, J. D., and Dutta, S., "Reconstruction of the Adaptable Deployable Entry and Placement Technology Sounding Rocket One Flight Test," *AIAA Aviation 2019 Forum*, AIAA, Dallas, TX, 2019. doi:10.2514/6.2019-2899.

[14] Leslie, F. W., and Justus, C. G., "The NASA Marshall Space Flight Center Earth Global Reference Atmosphere Model-2010 Version," *NASA/TM-2011-216467*, 2011.