MULTIFIDELITY APPROACHES FOR UNCERTAINTY ESTIMATION IN WILDFIRE SPREAD SIMULATORS

MARIO MIGUEL VALERO[1], LLUÍS JOFRE[2,3] AND RICARDO TORRES[3]

[1] Department of Meteorology and Climate Science, San José State University
San Jose, CA 95192, USA
mm.valero@sjsu.edu

[2] Center for Turbulence Research, Stanford University
Stanford, CA 94305, USA

[3] Department of Fluid Mechanics, Technical University of Catalonia - BarcelonaTech
Barcelona 08019, Spain
lluis.jofre@upc.edu; ricardo.torres@upc.edu

Key words: Computational fluid dynamics, Forest fire, Multifidelity modeling, Predictive science, Uncertainty quantification

Abstract. A variety of wildfire models are currently used for prescribed fire management, fire behaviour studies and decision support during wildfire emergencies, among other applications. All these applications are based on predictive analysis, and therefore require careful estimation of aleatoric and epistemic uncertainties such as weather conditions, vegetation properties and model parameters. However, the large computational cost of high-fidelity computaional fluid dynamics models prohibits the straightforward utilization of traditional Monte Carlo methods. Conversely, low-fidelity fire models are several orders of magnitude faster but they typically do not provide enough accuracy and they do not resolve all relevant phenomena. Multifidelity frameworks offer a viable solution to this limitation through the efficient combination of high- and low-fidelity simulations. While high-fidelity models provide the required level of accuracy, low-fidelity simulations are used to economically improve the confidence on estimated uncertainty. In this work, we assessed the suitability of multifidelity methodologies to quantify uncertainty in wildfire simulations. A collection of different multifidelity strategies, including Multilevel and Control Variates Monte Carlo, were tested and their computational efficiency compared. Fire spread was predicted in a canonical scenario using popular simulators such as the Wildland-Urban Interface Fire Dynamics Simulator (WFDS) and FARSITE. Results show that multifidelity estimators allow speedups in the order of $100\times$ to $1000\times$ with respect to traditional Monte Carlo.

1 INTRODUCTION

A variety of wildfire models have been developed with the aim of simulating forest fire behaviour and effects. Exiting models range from completely empirical correlations to detailed physics-based simulation frameworks [1, 2, 3]. Among other applications, these models are currently used for prescribed fire management, fire behaviour studies and decision support during wildfire emergencies. Tools designed to be used operationally for decision support need to be computationally fast and provide easy-to-interpret

information regarding macroscopic variables. Conversely, the detailed study of fire dynamics requires additional insight into the physical and chemical phenomena involved in fire spread, which can only be achieved by means of computational fluid dynamics (CFD). CFD approaches typically require intense computing resources that are several orders of magnitude larger than what operational simulators demand.

All these applications are based on predictive analysis, and therefore require careful estimation of the aleatoric and epistemic uncertainties, such as weather conditions, vegetation properties and model parameters. However, little to no information is available at present about wildfire simulation uncertainties. Although current simulators have been sporadically validated against experimental fire behaviour data, there has been no systematic analysis of their behaviour in terms of uncertainty propagation. One of the main reasons for this is the large computational cost of high-fidelity models, which prohibits the straightforward utilization of traditional Monte Carlo (MC) methods. Conversely, low-fidelity fire models are several orders of magnitude faster, but they typically do not provide enough accuracy.

Multifidelity (MF) frameworks offer a viable solution to this limitation through the efficient combination of high- and low-fidelity simulations. While high-fidelity (HF) models provide the required level of accuracy, low-fidelity (LF) simulations are used to economically improve the confidence on estimated uncertainty. Different MF UQ strategies exist in the literature; see, for example, the reviews by Peherstorfer et al. [4] and Fernández-Godino et al. [5, 6].

The objective of this work is to explore the potential of MF techniques to reduce the computing requirements of performing uncertainty quantification and sensitivity analysis in wildfire spread computational predictions. Due to the high-dimensional input space and the complexity of the conservation equations involved, this study is restricted to a reduced subset of acceleration strategies appertaining to surrogate-based MC type sampling approaches.

2 METHODOLOGY

In this study, we applied two state-of-the-art multifidelity modeling techniques to a canonical wildfire spread simulation problem. The required fidelity levels were built using two popular fire spread models of different typology: WFDS and FARSITE. HF simulations were produced using the Wildland-Urban Interface Fire Dynamics Simulator (WFDS) [7], a full-physics CFD model currently incorporated into the broader Fire Dynamics Simulator (FDS) [8, 9]. Increasingly coarser discretizations of the domain were used to generate lower fidelity levels from the HF scenario. Furthermore, an additional low fidelity representation of the problem was introduced through a semiempirical fire spread model. One of the most popular solutions among operational wildfire simulators is the Rothermel model [10]. Based on semi-empirical relationships between the parameters that determine the heat emitted by the fire and the energy needed by the unburned fuel to ignite, the Rothermel model provides a well-balanced combination of physical insight and operational capabilities. Rothermel's model constituted the base of FARSITE [11], a 2D fire spread simulator widely used for practical applications.

The analyzed scenario (Fig. 1) belongs to a set of medium-scale field fire experiments conducted by the Commonwealth Scientific and Industrial Research Organisation (CSIRO) in the Northern Territory of Australia in 1986 [12]. These tests were monitored to measure the fire rate of spread and evaluate its correlation with fuel and weather variables such as fuel height, fuel moisture content, fuel load, fuel bulk density, and wind speed. Two of the experiments had been simulated using WFDS [7] and incorporated

into the FDS test suite as validation examples [13]. Furthermore, the fact that the experiments were conducted on horizontal grassland fields facilitates the application of Rothermel's model for surface fire spread. In particular, this study considers the CSIRO grassland F19 experiment, which burned a square field of $200\,m \times 200\,m$ using a 175-meter long ignition line. This type of experiment design is very frequent in wildfire behavior studies.


Figure 1: Canonical wildfire spread scenario used to test MF modeling techniques. The CSIRO grassland F19 experiment [12, 7, 13] was reproduced using WFDS and FARSITE. This image is a snapshot outputted from WFDS for a sample run and displayed using Smokeview [14]. The yellow surface represents the field, the black dots correspond to the grass particles, the red line is the ignition line, and the orange region indicates the location of the fire as it spreads.

The two multifidelity techniques explored are Multilevel Monte Carlo and Control Variates Monte Carlo. As their name indicate, MC-type approaches are derived from the original Monte Carlo method, in which the expectation of the quantity of interest (QoI) as a function of the stochastic inputs ξ, $Q = Q(\xi)$, is estimated via a sample average. Let $\mathbb{E}[Q]$ and $\mathbb{V}[Q]$ denote the mean and variance of Q. Given N independent realizations of the stochastic input, denoted $\xi^{(i)}$, the MC estimator of $\mathbb{E}[Q]$ is defined as $\hat{Q}_N^{\mathrm{MC}} = N^{-1} \sum_{i=1}^{N} Q^{(i)}$, where $Q^{(i)} = Q(\xi^{(i)})$. Although unbiased, the precision of $\hat{Q}_N^{\mathrm{MC}}$, measured by its standard deviation $\sqrt{\mathbb{V}[Q]/N}$, decays slowly as a function of N. Therefore, for a fixed computational budget ($\propto N$), a viable alternative to increase the MC precision is to possibly replace Q with other quantities with smaller variances.

2.1 Multilevel Monte Carlo

The multilevel (ML) method [15, 16], inspired by the multigrid solver idea in linear algebra, is based on evaluating realizations of Q from a hierarchy of models with different levels ℓ, $\ell = 0,\ldots L$, with L the most accurate model, in which Q is replaced by the sum of differences $Y_\ell = Q_\ell - Q_{\ell-1}$; to simplify the notation for level 0, the expression is redefined to $Y_0 = Q_0$. As a result, the QoIs of the original and new ML problems have the same mean $\mathbb{E}[Q]$. An example of a level is the grid resolution considered for solving the system of equations, so that a LF (or HF) model can be established by simulating Q on a coarse (or fine) grid. Then, $\mathbb{E}[Q]$ can be computed using the ML QoI and an independent MC estimator on each level ℓ as

$$\hat{Q}^{\mathrm{ML}} = \sum_{\ell=0}^{L} \hat{Y}_\ell^{\mathrm{MC}} = \sum_{\ell=0}^{L} \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} Y_\ell^{(i)}. \tag{1}$$

This approach is referred to as multilevel Monte Carlo (MLMC), or simply ML, and the resulting estimator has a variance equal to

$$\mathbb{V}\left[\hat{Q}^{\mathrm{ML}}\right] = \sum_{\ell=0}^{L} \frac{\mathbb{V}\left[Y_\ell\right]}{N_\ell}. \tag{2}$$

Consequently, if the level definition is such that $Q_\ell \to Q$ in mean square sense, then $\mathbb{V}\left[Y_\ell\right] \to 0$ as $\ell \to \infty$, and therefore fewer samples are required on the finer level L. In particular, it is possible to show that the optimal sample allocation across levels N_ℓ is obtained in closed form given a target variance of the ML estimator equal to $\varepsilon^2/2$, and resulting in [15]

$$N_\ell = \frac{\sum_{k=0}^{L} \sqrt{C_k \mathbb{V}\left[Y_k\right]}}{\varepsilon^2/2} \sqrt{\frac{\mathbb{V}\left[Y_\ell\right]}{C_\ell}}, \tag{3}$$

where the computational cost of the individual Y_ℓ evaluations is denoted by C_ℓ, and ε^2 represents the mean squared error (MSE) of the estimator. It is important to note that the variance decay can be proven to be satisfied only for levels based on a numerical discretization (spatial/temporal meshes), and not for general hierarchies of models, such as 2-D versus 1-D, large-eddy simulation (LES) versus Reynolds-Averaged Navier Stokes (RANS), etc.

To apply MLMC, the HF CFD scenario was resolved using decreasing levels of resolution in (i) flow cell size, (ii) integration time-step (Δt), and (iii) number of discrete angles utilized to resolve the radiative heat transport. Based on preliminary tests, we designed three different levels named LF-2x, LF-5x and LF-10x, which are $25\times$, $250\times$ and $2500\times$ cheaper to compute than the HF level, respectively. The different resolution levels for (i) and (iii) are specified in Table 1. The integration time-step was automatically adjusted at every iteration based on the CFL constraint with CFL < 1; approximately, a cell-size increase of $2\times$ corresponds to a two-fold increase of Δt.

Table 1: Resolution levels used for MLMC

Level	Cell size (m)	# Discrete angles	HF-equivalent sample cost
HF	0.50	100	1
LF-2x	1.00	50	1/25
LF-5x	1.25	20	1/250
LF-10x	2.00	10	1/2500

2.2 Control variates Monte Carlo

To accommodate LF representations that are not obtained directly from coarsening the HF models, a common approach is to utilize LF realizations as a control variate [4, 17, 18]. In statistics, the control variates approach replaces a generic quantity q by $q + \alpha(\mathbb{E}\left[g\right] - g)$, where g is a function chosen for its high correlation with q and for which the value of $\mathbb{E}\left[g\right]$ is readily available. However, in the problem of interest here the LF correlations and expected values are not available *a priori*, and consequently need to be established during the computations along with the HF calculations. As a consequence, the expected values of the LF models are generally approximated by means of MC estimators requiring a set of additional (independent) LF computations. The control variates (CV) MC estimator is defined as

$$\hat{Q}^{\mathrm{CV}} = \hat{Q}^{\mathrm{MC}}_{\mathrm{HF}} + \alpha \left(\mathbb{E}\left[Q_{\mathrm{LF}}\right] - \hat{Q}^{\mathrm{MC}}_{\mathrm{LF}}\right), \tag{4}$$

where $\hat{Q}_{\mathrm{HF}}^{\mathrm{MC}} = N_{\mathrm{HF}}^{-1} \sum_{i=1}^{N_{\mathrm{HF}}} Q_{\mathrm{HF}}^{(i)}$, $\mathbb{E}[Q_{\mathrm{LF}}] \approx (N_{\mathrm{LF}} - N_{\mathrm{HF}} + 1)^{-1} \sum_{i=N_{\mathrm{HF}}+1}^{N_{\mathrm{LF}}} Q_{\mathrm{LF}}^{(i)}$, $\hat{Q}_{\mathrm{LF}}^{\mathrm{MC}} = N_{\mathrm{HF}}^{-1} \sum_{i=1}^{N_{\mathrm{HF}}} Q_{\mathrm{LF}}^{(i)}$, N_{HF} and N_{LF} are the number of HF and LF samples, respectively, and $\alpha = \mathbb{C}[Q_{\mathrm{HF}}, Q_{\mathrm{LF}}] / \mathbb{V}[Q_{\mathrm{LF}}]$ is the control variates coefficient chosen to minimize the variance of $\hat{Q}^{\mathrm{CV}}$. $\mathbb{C}[Q_{\mathrm{HF}}, Q_{\mathrm{LF}}]$ denotes the covariance between Q_{HF} and Q_{LF}. The optimal α selection leads to

$$\mathbb{V}\left[\hat{Q}^{\mathrm{CV}}\right] = \mathbb{V}\left[Q_{\mathrm{HF}}\right]\left(1 - \rho^2 \frac{r}{r+1}\right), \qquad (5)$$

with $-1 \le \rho = \mathbb{C}[Q_{\mathrm{HF}}, Q_{\mathrm{LF}}] / \sqrt{\mathbb{V}[Q_{\mathrm{HF}}] \mathbb{V}[Q_{\mathrm{LF}}]} \le 1$ the Pearson correlation coefficient between the HF and LF models, and r is used to parameterize the additional rN_{HF} LF realizations with respect to HF. As described by Geraci et al. [18], the optimal control variates is obtained for

$$r = \sqrt{\frac{C_{\mathrm{HF}}}{C_{\mathrm{LF}}} \frac{\rho^2}{1-\rho^2}} - 1, \qquad (6)$$

where C_{HF} and C_{LF} are the costs of a HF and LF sample, respectively. A comprehensive description of the control variates MC estimator, including the derivation of optimal coefficients and number of samples per fidelity, can be found in Peherstorfer et al. [4].

In this study, we introduced the output of the Rothermel fire spread model as a control variate. This was accomplished through the execution of FARSITE simulations in an scenario similar to the one displayed in Fig. 1. This fidelity is referred to as LF-CV, and is 2500× faster to compute than the HF level.

3 RESULTS

The performance of various candidate MF estimators constructed by means of CV and ML strategies was analyzed by utilizing a set of 32 pilot samples generated following a design of experiment (DoE) based on the KDOE approach [19, 20]. KDOE is an iterative method that introduces stochasticity in the sampling process by means of a variable kernel density estimation to optimize the uniformity of the DoE. This approach provides a more homogeneous exploration of the input parameter space, especially when the number of samples is relatively small. Five input variables were allowed to vary during pilot sampling: wind speed, fuel load, fuel bed depth, fuel moisture content and fuel particle surface area to volume ratio. Fire rate of spread (ROS), averaged over the complete simulation domain, was used as QoI for this demonstration.

As discussed in Section 2, the speedup obtained by the ML and CV approaches is function, respectively, of the variance of Y_ℓ ($\mathbb{V}[Y_\ell]$) in Eq. 2 and the Pearson correlation coefficient (ρ) between fidelities in Eq. 5. Consequently, $\mathbb{V}[Y_\ell]$ and ρ for all potential combinations are listed in Table 2. In the case of CV-based MF estimators, which are constructed utilizing HF information and LF samples as a control variate, the best LF candidates are LF-2x and LF-5x as they present correlations of $\rho = 0.995$ and $\rho = 0.990$ with speedups of approximately 25× and 250×, respectively. Instead, LF-CV and LF-10x are slightly less correlated with HF as their values are $\rho = 0.884$ and $\rho = 0.869$. If considering ML strategies, in which HF and different LF are combined forming a telescopic sum, a good hierarchical structure is composed by the lower fidelities LF-2x, LF-5x and LF-10x as their $\mathbb{V}[Y_\ell]$ values decay orders of magnitude, 0.002, 0.017 and 0.224 specifically, while becoming 25×, 250× and 2500× faster to compute than HF.

The extrapolated performances of a straightforward MC approach and the CV and ML estimators proposed above are reported in Fig. 2. The horizontal logarithmic-scale axis corresponds to the total cost of

Table 2: Pearson correlation coefficient ρ (elements below diagonal) and variance of levels $\mathbb{V}[Y_\ell]$ (elements above diagonal) for all the potential combinations of fidelities.

$\rho \setminus \mathbb{V}[Y_\ell]$	HF	LF-CV	LF-2x	LF-5x	LF-10x
HF	$1.0 \setminus 0.0$	0.303	0.002	0.024	0.312
LF-CV	0.884	$1.0 \setminus 0.0$	0.300	0.284	0.539
LF-2x	0.995	0.885	$1.0 \setminus 0.0$	0.017	0.291
LF-5x	0.990	0.889	0.993	$1.0 \setminus 0.0$	0.224
LF-10x	0.869	0.784	0.877	0.901	$1.0 \setminus 0.0$

each estimator normalized by the cost of a HF sample. The total costs are evaluated as $C^{\mathrm{MC}} = N_{\mathrm{HF}} C_{\mathrm{HF}}$, $C^{\mathrm{CV}} = N_{\mathrm{HF}} (C_{\mathrm{HF}} + r C_{\mathrm{LF}})$ and $C^{\mathrm{ML}} = \sum_{\ell=0}^{L} N_\ell C_\ell$ for the MC, CV and ML, respectively. On the vertical logarithmic-scale axis, target estimators' MSE, $\varepsilon^2 \equiv \mathbb{V}[\hat{Q}]$, normalized by a reference MC value $\varepsilon_0^{\mathrm{MC}2}$ obtained from the 32 pilot samples, are shown for MC and ML estimators through evaluating Eq. 3, and for the CV estimators utilizing the expression [4]

$$\frac{\varepsilon^2}{\varepsilon_0^{\mathrm{MC}2}} = \left[\sqrt{1-\rho^2} + \sqrt{\frac{C_{\mathrm{LF}}}{C_{\mathrm{HF}}} \rho^2} \right]^2. \tag{7}$$

The results depicted in Fig. 2 show that the speedups of the MF estimators with respect to MC are in the order of $100\times$ to $1000\times$ for the best-performant CV (constructed with HF and LF-5x) and ML (generated by levels Y_0: LF-10x, Y_1: LF-5x $-$ LF-10x, Y_2: LF-2x $-$ LF-5x and Y_3: HF $-$ LF-2x) strategies, respectively. For example, to reduce ε^2 by an order of magnitude with respect to $\varepsilon_0^{\mathrm{MC}2}$, the MC approach requires computing 320 HF runs, while the CV (composed by 1 HF and 493 LF-5x samples) and ML (composed by 1 HF, 19 LF-2x, 222 LF-5x and 1556 LF-10x samples) demand only 3.0 and 2.8 equivalent HF runs.

The results shown in Fig. 2 also highlight the better performance of the ML with respect to the CV estimators in the case of LF models/levels that present small bias and moderate CoV. However, in a more general problem involving very complex, non-linear fire spreads, in which such "good" LF models are more challenging to design and/or discover, the CV approach may be a more robust option. Thus, hybridization strategies, like for example the bi-fidelity (BF) approximation [21, 22, 23, 24] and the multilevel multifidelity (MLMF) approach [18, 25], are promising extensions of the standard ML and CV methods for accelerating the convergence rate of statistical estimators in challenging wildfire spread scenarios.

4 CONCLUSIONS

Performing uncertainty quantification in wildfire modeling studies is usually challenging due to the expensive high-fidelity calculations required and the large number of uncertainties typically encountered. By constructing and using multifidelity estimators, we were able to notably accelerate the propagation of aleatoric uncertainty through a CFD framework.

Multilevel Monte Carlo performed slightly better than Control Variates Monte Carlo due to the small bias and moderate variability of the low-fidelity data generated. However, in a more general problem involving very complex fire behavior, in which "good" low-fidelity models are challenging to design
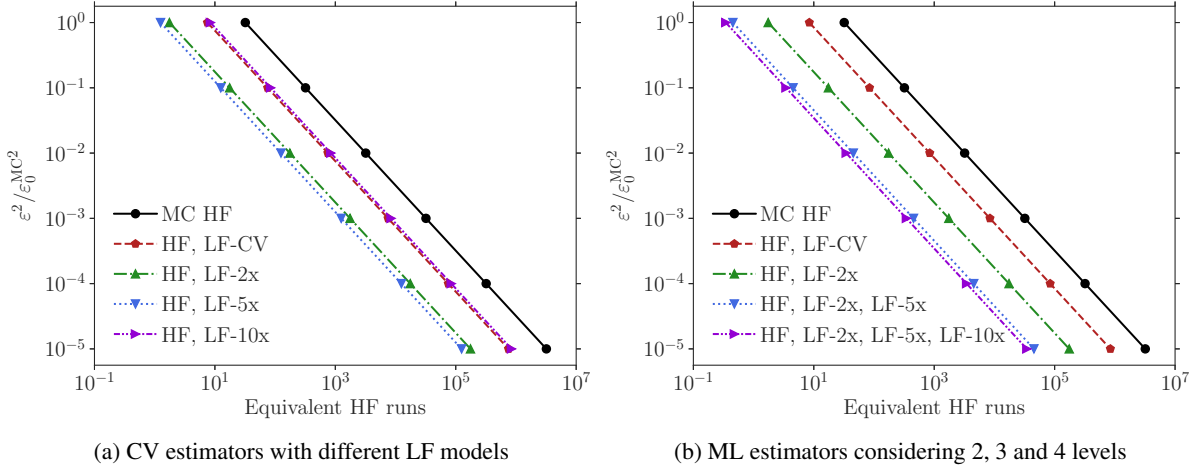
(a) CV estimators with different LF models

(b) ML estimators considering 2, 3 and 4 levels

Figure 2: Extrapolated MSE (normalized by the pilot $\varepsilon_0^{MC^2}$ value) of the MC and potential combinations of MF estimators as function of the overall computational cost in terms of equivalent number of HF runs. Solid black lines correspond to plain MC with HF samples.

and/or discover, the control variates approach may be a more robust option.

This proof of concept opens an avenue of further research. Future work includes the quantification of aleatoric and epistemic uncertainties, the study of multiple quantities of interest, and the analysis of model sensitivity to input variables as well as internal parameters.

ACKNOWLEDGEMENTS

References

[1] A. L. Sullivan. Wildland surface fire spread modelling, 1990–2007. 1: Physical and quasi-physical models. *International Journal of Wildland Fire*, 18:349–368, 2009.

[2] A. L. Sullivan. Wildland surface fire spread modelling, 1990–2007. 2: Empirical and quasi-empirical models. *International Journal of Wildland Fire*, 18:369–386, 2009.

[3] A. L. Sullivan. Wildland surface fire spread modelling, 1990–2007. 3: Simulation and mathematical analogue models. *International Journal of Wildland Fire*, 18:387–403, 2009.

[4] B. Peherstorfer, K. Willcox, and M. Gunzburger. Survey of multifidelity methods in uncertainty

propagation, inference, and optimization. *SIAM Rev.*, 60(3):550–591, 2018.

[5] M. G. Fernández-Godino, C. Park, N. H. Kim, and R. T. Haftka. Review of multi-fidelity models. *arXiv preprint arXiv:1609.07196*, 2016.

[6] M. G. Fernández-Godino, C. Park, N. H. Kim, and R. T. Haftka. Issues in deciding whether to use multifidelity surrogates. *AIAA Journal*, 57:2039–2054, 2019.

[7] W. Mell, M.A. Jenkins, J. Gould, and P. Cheney. A physics based approach to modeling grassland fires. *Int. J. Wildland Fire*, 16(201):1–22, 2007.

[8] Glenn P. Forney, Daniel Madrzykowski, Kevin B. McGrattan, and Laurel Sheppard. Understanding fire and smoke flow through modeling and visualization. *IEEE Computer Graphics and Applications*, 23(4):6, 2003.

[9] Jason Floyd and Kevin McGrattan. Validation of a CFD fire model using two step combustion chemistry using the NIST reduced-scale ventilation-limited compartment data. In *Fire Safety Science-Proceedings of the Ninth International Symposium*, pages 117–128, 2008.

[10] Richard C Rothermel. A mathematical model for predicting fire spread in wildland fuels. *USDA Forest Service Research Paper INT USA*, (INT-115):40, 1972.

[11] Mark A. Finney. FARSITE : Fire Area Simulator — Model Development and Evaluation. Technical report, U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station, 1998.

[12] N. P. Cheney, J. S. Gould, and W. R. Catchpole. The Influence Of Fuel, Weather And Fire Shape Variables On Fire-Spread In Grasslands. *International Journal of Wildland Fire*, 3(1):31–44, 1993.

[13] Kevin B. McGrattan, Randall J. McDermott, Marcos Vanella, Simo Hostikka, and Jason Floyd. Fire Dynamics Simulator Technical Reference Guide Volume 3: Validation, 2019.

[14] Glenn P. Forney. Smokeview, A Tool for Visualizing Fire Dynamics Simulation Data. Volume I: User's Guide, 2016.

[15] M. B. Giles. Multi-level Monte Carlo path simulation. *Oper. Res.*, 56:607–617, 2008.

[16] C. Adcock, Y. Ye, L. Jofre, and G. Iaccarino. Multilevel Monte Carlo sampling on heterogeneous computer architectures. *Int. J. Uncertain. Quan.*, 10:575–594, 2020.

[17] R. Pasupathy, M. Taaffe, B. W. Schmeiser, and W. Wang. Control-variate estimation using estimated control means. *IIE Trans.*, 44:381–385, 2014.

[18] G. Geraci, M. Eldred, and G. Iaccarino. A multifidelity multilevel Monte Carlo method for uncertainty propagation in aerospace applications. *19th AIAA Non-Deterministic Approaches Conference*, pages AIAA 2017–1951, 2017.

[19] P. T. Roy, L. Jofre, J.-C. Jouhaud, and B. Cuenot. Versatile sequential sampling algorithm using Kernel Density Estimation. *Eur. J. Oper. Res.*, 284:201–211, 2020.

[20] M. Masquelet, J. Yan, A. Dord, G. Laskowski, Lee Shunn, L. Jofre, and G. Iaccarino. Uncertainty quantification in large eddy simulations of a rich-dome aviation gas turbine. In *Proc. ASME Turbo Expo 2017, GT2017-64835*, pages 1–11, 2017.

[21] H. R. Fairbanks, A. Doostan, C. Ketelsen, and G. Iaccarino. A low-rank control variate for multi-

level Monte Carlo simulation of high-dimensional uncertain systems. *J. Comput. Phys.*, 341:121–139, 2017.

[22] L. Jofre, G. Geraci, H. R. Fairbanks, A. Doostan, and G. Iaccarino. Multi-fidelity uncertainty quantification of irradiated particle-laden turbulence. *CTR Annu. Res. Briefs*, pages 21–34, 2017.

[23] H. R. Fairbanks, L. Jofre, G. Geraci, G. Iaccarino, and A. Doostan. Bi-fidelity approximation for uncertainty quantification and sensitivity analysis of irradiated particle-laden turbulence. *J. Comput. Phys.*, 402:108996, 2020.

[24] L. Jofre, M. Papadakis, P. T. Roy, A. Aiken, and G. Iaccarino. Multifidelity modeling of irradiated particle-laden turbulence subject to uncertainty. *Int. J. Uncertain. Quan.*, 10:499–514, 2020.

[25] A. A. Gorodetsky, G. Geraci, M. S. Eldred, and J. D. Jakeman. A generalized approximate control variate framework for multifidelity uncertainty quantification. *arXiv preprint*, arXiv:1811.04988:[stat.CO], 2019.