

ARTICLE TYPE

An Adjoint-assisted Multilevel Multifidelity Method for Uncertainty Quantification And Its Application To Turbomachinery Manufacturing Variability

Pavanakumar Mohanamuraly*¹ | Jens Dominik Müller²

¹ALGO-COOP Team, CERFACS, Toulouse, France

²School of Engineering and Material Science, Queen Mary University of London, London, United Kingdom

Correspondence

*Corresponding author Email: mpkumar@cerfacs.fr

Summary

In this work we propose, analyse, and demonstrate an adjoint-based multilevel multifidelity Monte Carlo (MLMF) framework called FastUQ. The framework is based on the MLMF of Geraci et al.¹ and uses the Inexpensive Monte Carlo (IMC) method of Ghate² as low-fidelity surrogate. The setup cost of IMC-1 surrogate in FastUQ requires just the adjoint solution at the input mean whose computational cost is independent of the number of input uncertainties making it suitable for solving problems with a large number of uncertain parameters. We demonstrate the robustness of the method to quantify uncertainties in aerodynamic parameters due to surface variations caused by the manufacturing processes for a highly loaded turbine cascade. A stochastic model for surface variations on the cascade is proposed and optimal dimensionality reduction of model parameters is realised using goal-based principal component analysis using adjoint sensitivities of multiple quantities of interest. The proposed method achieves a 70% reduction in computational cost in predicting the mean quantities such as total-pressure loss and mass flow rate compared to the state-of-art MLMC method.

KEYWORDS:

Uncertainty Quantification, Multilevel Multifidelity Monte Carlo, Goal-based PCA, Adjoint Sensitivity, Manufacturing Variations

1 | INTRODUCTION

Determining performance deterioration of turbomachines due to manufacturing variations is important to engine manufacturers when determining trade-offs between manufacturing cost and performance. Many studies have been conducted in the literature with focus on change in performance due to manufacturing variations on blade shapes^{3,4,5,6,7,8,9}. The main challenge of modelling the effect of geometric variations on the blade performance is the large number of uncertain parameters that have to be considered. Lange et al.⁸ considered blade measurements from 3-D surface scan data of a high pressure compressor stage. The correlated coordinates from the 3-D scan were translated into mode shapes and uncorrelated modal amplitudes using PCA. Lange et al. included all 135 PCA modes in their perturbation model within the MC simulations and compared it with MC results considering 20 and 60 PCA modes. Lange⁸ concludes that 60 PCA modes (half of the total number of PCA modes) were necessary to correctly model the manufacturing variability on the blades. Even with dimensionality reduction the resulting parameter space

is too high⁸ to achieve the specified error tolerance. Therefore uncertainty quantification (UQ) methods that suffer from the curse of dimensionality (CoD) are impractical to solve this problem¹⁰.

UQ using the Monte Carlo method (MC) has been widely adopted^{8,9,4,5} due to the following key reasons; (i) convergence of statistical moments is independent of the number of input uncertainties^{11,12}, (ii) computational cost versus statistical errors can be established clearly hence the statistical estimates can be improved on the fly with additional samples¹³ and (iii) the method can be applied to highly non-linear and non-smooth systems¹⁴. The computational cost of MC becomes prohibitively expensive for problems where the computational cost per sample is high due to the slow convergence with the number of samples. This has motivated many improvements to the method, specifically the control variate acceleration approach¹.

The control variate naturally lends itself to multifidelity and multilevel formulations. In the context of robust aerospace design, Ng and Wilcox¹⁵ proposed a practical multifidelity control variate Monte Carlo (MFMC) implementation and its extension to an information reuse estimator¹⁶ that incorporates information from previous design steps. The adaptive MFMC (AMFMC) of Peherstorfer et al.¹⁷ splits the computational budget between adapting the low-fidelity models to improve their approximation quality and sampling the low- and high-fidelity models to reduce the mean square error of the estimator. Giles¹⁸ proposed the Multilevel Monte Carlo (MLMC) method, which uses a sequence of coarsened meshes to represent the low fidelity model of the control variate. Mishra¹⁹ et al. showed that MLMC has the same asymptotic complexity (up to a logarithmic term) as a single deterministic solution considering a hierarchy of nested meshes. Geraci et al.¹ proposed a multilevel multifidelity framework (MLMF) to achieve further reduction in computational cost by combining the multifidelity control variate in MLMC. Gorodetsky²⁰ gives a comprehensive overview of MLMF models and unify many existing sampling methods used for MF UQ. They propose non-recursive/nested strategies for sampling schemes and achieve significant computational gains.

The MFMC, AMFMC or MLMF frameworks require a surrogate model for the low-fidelity sample evaluation. Surrogate models do incur setup cost through initial high-fidelity sampling thus suffering from CoD. For example in the Kriging meta-model requires design of experiments (DOE) to sample the response at multiple points in the parameter space to create the Kriging model. Mostly the Latin hypercube sampling (LHS) is used for sampling the response for DOE^{14,21}. Nobile²² used a stochastic collocation to construct the low-fidelity surrogate model at each level in the MLMC for elliptic type PDE. Reynolds averaged Navier-Stokes (RANS) equations considered in this work is a highly non-linear mixed PDE that admits various types of flow discontinuities. In addition, to construct the surrogate model one requires samples at the sparse grid collocation points. Hence all the aforementioned methods are not suitable for problems with high input dimension.

Fairbanks²³ introduced a low-rank approximation in MLMC for higher dimensional problems called the multilevel control variates (MLCV). The low-rank approximation is obtained using a small set of selected samples of the fine grid solution identified from the realisations of the coarse grid solution. The method outperforms MLMC when the problem exhibits small solution ranks on all the levels. In addition, they show that MLCV is beneficial over MLMC only for low mean-squared error tolerance. Blonigan²⁴ used a Reduced Order Model (ROM) in a MF modelling approach based on the MF model of Gorodetsky²⁰ to solve the Kuramoto-Sivashinsky equation.

Ghate² proposed an adjoint-based surrogate model called the incomplete Monte Carlo (IMC) as a replacement for high fidelity flow simulations to study the performance of turbomachines and airfoils subject to random manufacturing variations. A similar model was used by Rumpfkeil^{25,26} as an approximate surrogate model for lift and drag cost functions of an airfoil in aerodynamic shape design. Engels-Putzka et al.²⁷ used the IMC model to estimate the mass flow rate, total-pressure ratio and total-temperature ratio in the two-stage Darmstadt Transonic Compressor using the TRACE flow and adjoint solver developed at DLR for use in robust design. Luo et al.²⁸ adopted a similar approach of using the adjoint and higher order sensitivities to approximately evaluate the aerodynamic performance of a turbine subject to manufacturing variability using Taylor expansion. They used a continuous adjoint approach based on inviscid Euler equations to obtain the adjoint sensitivity and Hessian solutions. The adjoint gradient used in the IMC model of Ghate has the property that its computational cost is proportional to the number of output variables and independent of input dimension. Therefore the adjoint surrogate combined with MFMC or MFML forms an attractive combination because the convergence of the resulting method is still independent of the input dimension.

We find that in all the previous works based on the adjoint-based surrogates (similar to Ghate's IMC), the model was used as a replacement for the high-fidelity evaluation. We demonstrate in this work the weakness of using the IMC surrogate model for highly non-linear problems using the uncertain viscous Burgers' equation and aerodynamic performance variations due to manufacturing uncertainties on the surface of a turbine blade. Errors as large as 35 – 225% in standard deviation and 2 – 11% in mean values were obtained for the viscous Burgers' problem and 6% variation in mean and 10% variation in standard deviation was observed for the total-pressure loss for the turbine problem using IMC. In fact the higher-order adjoint corrections have larger errors compared to the lower order ones. Interestingly higher and lower order IMC models correlate well with the exact

solution. Therefore we propose to use the adjoint-surrogate as a low-fidelity model in the proposed multifidelity multilevel UQ framework. To the authors' knowledge this is the first reported work on the use of the adjoint-based IMC surrogate as a low-fidelity model within a MLMF framework.

The MLMF framework proposed in this work is based on the work of Geraci et al.¹ but we use the multifidelity cost model similar to the one used in Ng¹⁵. This clearly brings out the definition of the multifidelity parameter in the multilevel multifidelity framework. We introduce a Gaussian process model for surface variations due to manufacturing process of turbine blades. Unlike previous work in the literature^{29,27,4,6} we consider multiple cost functions in a goal-based principal component analysis of the surface variation model. We show that including multiple cost functions increases the the number of PCA modes required to achieve the specified error tolerance in the cost functions. This corroborates with the high dimensional nature of the problem of uncertainty in aerodynamic performance due to manufacturing variations of blades. Lastly to the authors knowledge this one of the first work to use an MLMF framework for the problem of aerodynamic UQ due manufacturing variations.

2 | PROBLEM FORMULATION

2.1 | Multifidelity control variate acceleration

Let J be a quantity of interest (QoI), which we try to estimate using P high-fidelity (HF) samples to yield $J_p : \{p = 1, \dots, P\}$. The mean-squared error (MSE) in QoI J using the classical Monte Carlo (MC)³⁰ on P samples is a sum of the bias error and the stational error as shown in Eq. (1). $\mathbb{E}[\cdot]$ and $\text{Var}[\cdot]$ denote the mean and variance.

$$\text{MSE}[J_p^{MC}] = \mathbb{E} \left[\left(J_p^{MC} - \mathbb{E}[J] \right)^2 \right] = \underbrace{P^{-1} \text{Var}[J_p]}_{\text{Statistical error}} + \underbrace{\left(\mathbb{E}[J_p] - \mathbb{E}[J] \right)}_{\text{Bias error}} \quad (1)$$

The term J_p^{MC} is the MC estimate of the mean obtained using the P samples i.e., $J_p^{MC} = P^{-1} \sum_{p=1}^P J_p$. For an unbiased estimator the MSE in eq.(1) can be reduced either by, (1) increasing the number of samples P or (2) reducing the variance of J_p . The control variate takes the latter approach and reduces J_p by using M samples from a low-fidelity (LF) model G ($G_m : \{m = 1, \dots, M\}$). Typically G is much cheaper to evaluate than J although less accurate. Variance reduction is achieved by introducing the new estimator involving J , G and a free parameter β (called the control variate parameter) shown in Eq. (2).

$$J_p^{CV} = J_p^{MC} - \beta (G_m^{MC} - \mathbb{E}[G]) \quad (2)$$

G_m^{MC} is the MC estimate of the mean value of G using M samples and $\mathbb{E}[G]$ is the actual mean. The control variate parameter β is obtained by minimising the variance of J_p^{CV} i.e., $\frac{d}{d\beta} (\text{Var}[J_p^{CV}]) = 0$, considering a suitable cost model for the sample evaluation. We use a cost model (similar to Pasupathy³¹ and Ng¹⁶) for the computational budget shown in Eq. (3) to minimise the variance.

$$c = P + \frac{M}{w} = P \left(1 + \frac{r}{w} \right) \quad (3)$$

In Eq. (3), w is the runtime ratio between the average computational cost per HF and LF model evaluation and $r = \frac{M}{P} > 1$ is the ratio between the number of HF and LF simulations. c is the computational budget (an equivalent HF sample count). For a fixed budget c we are free to choose the parameter r to vary the resource allocation between HF and LF evaluations. The variance of the multifidelity estimator in terms of c , r , and w is shown in Eq. 4. Pasupathy et al³¹ propose a novel control variate called Control Variates using Estimated Means (CVEM) where the authors introduce an additional source of sampling variability modelled using an analysis error $\varepsilon = \tilde{G}_m^{MC} - \mathbb{E}[G_m^{MC}]$. Where \tilde{G}_m^{MC} is an estimate of the approximation-model mean $\mathbb{E}[G_m^{MC}]$. Pisoni³² uses a weighted formulation of bias and statistical error in a multilevel control variate formulation called the continuation MLMC. In this work we consider the analysis error ε to be negligibly small (similar to the work of Ng¹⁶ and Geraci¹). This assumption is justified when the number of LF samples M used to estimate the mean $\mathbb{E}[G]$ is much greater than the number of high fidelity samples P , i.e. ($M \gg P$).

$$\begin{aligned} \text{Var}[J_p^{CV}] &= \frac{1}{c} \left(1 + \frac{r}{w} \right) \left[\text{Var}[J] + \left(1 - \frac{1}{r} \right) \left(\beta^2 \text{Var}[G] - 2\beta\rho\sqrt{\text{Var}[G]\text{Var}[J]} \right) \right] \\ &= \left[1 + \left(1 - \frac{1}{r} \right) \left(\beta^2 \frac{\text{Var}[G]}{\text{Var}[J]} - 2\beta\rho\sqrt{\frac{\text{Var}[G]}{\text{Var}[J]}} \right) \right] \frac{\text{Var}[J]}{P}. \end{aligned} \quad (4)$$

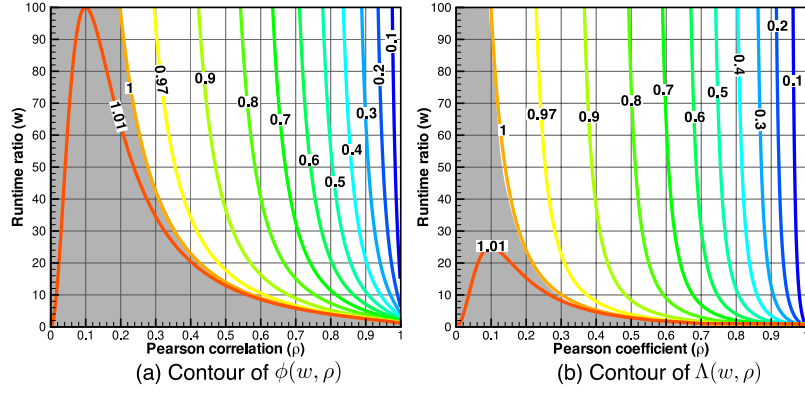


FIGURE 1 Contour plot of $\phi(r^{opt})$ and $\Lambda(r^{opt})$ for various values of runtime ratio w and Pearson correlation ρ (region where $\phi(r^{opt}) \geq 1$ and $\Lambda(r^{opt}) \geq 1$ is indicated in grey)

For optimal resource allocation, one minimises the variance in Eq. 4 for both β and r to yield the optimality condition¹⁵ in Eq. 5.

$$\text{MSE}[J_p]^{CV} = \text{Var}[J_p]_{opt}^{CV} = \left(1 + \frac{r_{opt}}{w}\right) \left(1 - \frac{r_{opt} - 1}{r_{opt}} \rho^2\right) \frac{\text{Var}[J]}{c} = \left(1 - \frac{r_{opt} - 1}{r_{opt}} \rho^2\right) \frac{\text{Var}[J]}{P} \quad (5)$$

where, r_{opt} and β_{opt} are,

$$r_{opt} = \sqrt{\frac{w\rho^2}{1 - \rho^2}}, \quad \text{and} \quad \beta_{opt} = \rho \sqrt{\frac{\text{Var}[J]}{\text{Var}[G]}} \quad (6)$$

When the cost of computing the LF samples is almost free i.e., $w \rightarrow \infty$, then $r_{opt} \rightarrow \infty$ and one recovers the classical control variate estimator i.e., $[J_p]^{CV} = (1 - \rho) \frac{\text{Var}[J]}{c}$. In the case of a perfect LF model i.e., $\rho \rightarrow 1$ then the error scales similar to the statistical error of a classical MC considering only the LF model. Therefore, the LF model should not only correlate well but also should be cheaper to evaluate.

2.2 | Convergence comparison of single- and multi-fidelity MC

A comparison with a classical MC estimator can be obtained by observing that the term $c^{-1} \text{Var}[J]$ is actually the MSE of MC estimator considering the same budget c (neglecting the bias error), i.e., $\text{MSE}[J_c^{MC}]$. Then dividing eq (5) by $\text{MSE}[J_c^{MC}]$ we obtain the ratio ϕ between the MSE of the optimal multifidelity and classical MC estimator under the same computational budget c shown in Eq. (7).

$$\phi = \frac{\text{MSE}[J_p]^{CV}}{\text{MSE}[J_c]^{MC}} = \left(1 + \frac{r_{opt}}{w}\right) \left(1 - \frac{r_{opt} - 1}{r_{opt}} \rho^2\right) \quad (7)$$

A similar comparison of the MF estimator can be done with the MSE of a classical MC using P samples from Eq. 6 (neglecting bias error). The parameter $\Lambda_l^{1,20}$ is a penalising factor with respect to an optimal control variate, which is the control variate obtained by using a low-fidelity model for which the statistics are known and therefore the variance reduction is $(1 - \rho^2)$. The contour plots of ϕ and Λ is plotted in Fig. 1. The grey regions denote the regions where the control variate is ineffective and a classical MC performs better.

$$\Lambda = \frac{\text{MSE}[J_p]^{CV}}{\text{MSE}[J_p]^{MC}} = \left(1 - \frac{r_{opt} - 1}{r_{opt}} \rho^2\right) \quad (8)$$

Note that there are no terms involving P in Eq. (7) so the multifidelity MC has the same $\mathcal{O}(P^{-\frac{1}{2}})$ convergence but the entire convergence curve is shifted by Λ . The shift is only a function of the correlation ρ and runtime ratio w .

2.3 | Multilevel Multifidelity Control Variate

Evaluation of the computational model on a coarser mesh reduces computational cost and accuracy. Using sample evaluations on a hierarchy of coarser meshes as LF model in the control variate method produces the multilevel MC (MLMC)¹⁸. Let J_N be the QoI obtained from solving a PDE on a mesh with N degrees-of-freedom. Therefore as $N \rightarrow \infty$ then $J_N \rightarrow J$ and the mean values $\mathbb{E}[J_N] \rightarrow \mathbb{E}[J]$. Using the linearity property of expectation, the expectation at the finest level L (with N_L degrees of freedom) can be written using the telescopic sum of the sequence of levels $l = 0, \dots, L$ as shown in Eq. (9). The difference function Y_l defined in Eq. (10) is introduced to simplify the analysis.

$$\mathbb{E}[J_N] = \mathbb{E}[J_{N_0}] + \sum_{l=0}^L \mathbb{E}[J_{N_l} - J_{N_{l-1}}] = \sum_{l=0}^L \mathbb{E}[Y_l] \quad (9)$$

$$Y_l = \begin{cases} J_{N_0} & \text{if } l = 0 \\ J_{N_l} - J_{N_{l-1}} & \text{if } l \neq 0 \end{cases} \quad (10)$$

For P_l samples distributed per level we get the MLMC estimator in Eq. (11) of Giles¹⁸ who also derived the optimal samples per level P_l^{opt} by optimising the total computational cost $c = \sum_{l=0}^L P_l C_l$ for a specified MSE tolerance ϵ on the estimator in Eq. (11) to yield Eq. (12). C_l is the cost of a single sample evaluation at level l . Note that Giles bounds the bias error in estimating Y_l to be less than the prescribed tolerance divided by the square root of two ($\frac{1}{\sqrt{2}}\epsilon$).

$$J_{N,p}^{ML} = \sum_{l=0}^L \frac{1}{P_l} \sum_{i=1}^{P_l} \mathbb{E}[Y_l^{(i)}] \quad (11)$$

$$P_l^{opt} = \frac{2}{\epsilon^2} \left[\sum_{k=0}^L \sqrt{\text{Var}(Y_k) C_k} \right] \sqrt{\frac{\text{Var}(Y_l)}{C_l}} \quad (12)$$

Mishra¹⁹ et al. showed that MLMC has the same asymptotic complexity (up to a logarithmic term) as a single deterministic solution considering a hierarchy of nested meshes. Therefore MLMC can converge QoI statistics rapidly; analogous to the multigrid convergence for PDE solution. Further variance reduction can be achieved by combining the multifidelity control variate into a multilevel one. Geraci et al.¹ combined both MFMC and MLMC and proposed a multilevel multifidelity MC (MLMF) framework. The number of LF samples used per level in the work of Geraci was fixed as an increment Δ of the number of HF samples $M = P + \Delta$, where $\Delta = \kappa P$ is some fraction of the HF samples. The multifidelity computational budget per level of Geraci using our notation is,

$$c = P \left(1 + \frac{1 + \kappa}{w} \right) \quad (13)$$

Comparing the multifidelity cost model in Eq. (13) with the one in Eq. (3) we obtain the equation $r = 1 + \kappa$, which relates the parameters used in both models. Replacing the cost model gives us the benefit of using the multifidelity estimator ϕ and clearly show the effect of introducing the multilevel into the multifidelity control variate, which is shown in detail in Sec. ???. To simplify the analysis the same number of grid levels are used for models J and G . The estimators in Eq. (2) and Eq. (11) can be combined to obtain the Y_l correlation MLMF estimator¹³, $J_{N,p}^{MLMF}$ of QoI J shown in Eq. (14).

$$J_{N,p}^{MLMF} = \sum_{l=0}^L \frac{1}{P_l} \sum_{i=1}^{P_l} \mathbb{E}[\bar{Y}_l^{(i)}] \quad (14)$$

where, \bar{Y}_l is the control variate at level l shown in Eq. (15) and Z_l is the difference function¹ of the LF model G shown in Eq. (16).

$$\bar{Y}_l = Y_l - \beta_l (Z_l - \mathbb{E}[Z_l]) \quad (15)$$

$$Z_l = \begin{cases} G_{N_0} & \text{if } l = 0 \\ G_{N_l} - G_{N_{l-1}} & \text{if } l \neq 0 \end{cases} \quad (16)$$

2.4 | Optimal sample allocation between levels and models

Let P_l and M_l denote the number of sample evaluations of the HF and LF models at level l . Let r_l denote the fraction of HF evaluation to the LF evaluation $r_l = M_l/P_l$. The cost of computation on a level is expressed in terms of the equivalent number

of HF sample evaluations as shown in Eq. (17), where C_l is a measure of the computation cost of HF sample on level l . We normalise C_l by dividing it by the computational cost of evaluating a HF sample at the finest level N (note that $C_N = 1$ and $C_l \leq 1$). Therefore the budget c_l now contains both the relative computational cost between models and across levels,

$$c_l = C_l P_l \left(1 + \frac{r_l}{w_l} \right). \quad (17)$$

The total computational budget c considering all levels $l = 0, \dots, L$ is then

$$c = \sum_{l=0}^L c_l = \sum_{l=0}^L C_l P_l \left(1 + \frac{r_l}{w_l} \right) \quad (18)$$

The variance of this estimator can be optimised for β_l and r_l using a Lagrangian multiplier approach¹. The final optimal variance and the optimal number of HF samples P_l^{opt} per level is shown in Eq. (19)-(20).

$$\text{Var} \left[J_{N, p}^{MLMF} \right] = \sum_{l=0}^L \frac{1}{P_l} \text{Var}[Y_l] \Lambda(r_l), \quad (19)$$

$$P_l^{opt} = \frac{2}{\epsilon^2} \left[\sum_{k=0}^L \sqrt{\frac{\text{Var}(Y_k) C_k}{1 - \rho_k^2}} \Lambda_k(r_k^{opt}) \right] \sqrt{(1 - \rho_l^2) \frac{\text{Var}(Y_l)}{C_l}}. \quad (20)$$

The values of r_l^{opt} and Λ_l are shown in Eq. (21)-(22) respectively. The parameter $\Lambda_l^{1,20}$ is a penalising factor with respect to an optimal control variate, which is the control variate obtained by using a low-fidelity model for which the statistics are known and therefore the variance reduction is $(1 - \rho^2)$. ρ_l is the Pearson correlation between the HF and LF model.

$$r_l^{opt} = \sqrt{\frac{\rho_l^2}{1 - \rho_l^2} w_l} \quad (21)$$

$$\Lambda_l(r_l) = \left(1 - \frac{r_l - 1}{r_l} \rho_l^2 \right) \quad (22)$$

To simplify the analysis, consider the correlation between the HF and LF model across levels to be the same ($\rho_l = \rho$). Then the optimal number of samples in Eq. (20) reduces to the simple expression in Eq. (23).

$$P_l^{opt} = \frac{2}{\epsilon^2} \sum_{k=0}^L \left[\sqrt{\text{Var}(Y_k) c_k} \Lambda_k(r_k^{opt}) \right] \sqrt{\frac{\text{Var}(Y_l)}{c_l}} \quad (23)$$

The optimal number of samples per level P_l^{opt} in Eq. (23) is identical to the one obtained by Giles (shown in Eq. (12)) except for the additional Λ term in the summation. In fact, for the special case of $\Lambda = 1$ we recover Eq. (12). The contour plot comparing $\phi(r^{opt})$ and $\Lambda(r^{opt})$ for a range of values of correlation ρ and runtime ratio w is shown in Fig. 1a)-b). We find that the runtime ratio w should be ≥ 10 and correlation ρ should be ≥ 0.6 for the MLMF to be effective. P_l^{opt} is directly proportional to $\Lambda_l(r_l^{opt})$, which is a strong function of the correlation. For large runtime ratios w , $\Lambda(r^{opt}) \propto \rho$ and is almost independent of runtime ratio. Therefore the success of the MLMF is strongly dependent on how well the LF model correlates with the HF model. In addition $P_l^{opt} \propto \epsilon^{-2}$ so lowering the tolerance by half quadruples the number of samples per level similar to the convergence of MC. The value of $\Lambda(r^{opt})$ is mostly ≤ 1 indicating a reduction in the number of samples compared to the standard MLMC to achieve the same MSE.

2.5 | Note on pilot samples

Practically, in the multifidelity control variate it is not possible to compute $\mathbb{E}[J_p]^{CV}$ and $\text{Var}[J_p]^{CV}$ using Eq. 4 and Eq. 5, because $\mathbb{E}[J]$, $\text{Var}[J]$, $\text{Var}[G]$, ρ and β are unknown quantities and the usual practice is to replace them with sample estimates, which are called pilot samples. The pilot sampling is also necessary for evaluating the variance estimate at a given level $\text{Var}(Y_k)$. For a budget c , one has N samples of J_p , which can be used to estimate the necessary statistics. Replacing the exact values with approximate ones introduces errors. Therefore, it is necessary to check if the estimator still yields a reduction in variance. This is checked by plotting the ratio ψ shown in Eq. 24 (from ref.¹⁵). In Fig. 2 the sample contours for ψ is shown for different values of r and β for two values of correlation ρ (a) 0.5 and (b) 0.75. For small errors in the sample estimate of β and r (exact values




FIGURE 2 Contours of ψ for indicated ratio of number of model evaluations r and control parameter β for (a) $\rho = 0.5$ and (b) $\rho = 0.75$.

indicated by \times) are acceptable but optimal variance reduction cannot be achieved¹⁵.

$$\psi = \frac{\mathbb{V}ar[J_p]^{CV}}{\mathbb{V}ar[J_p]} = \left(1 + \frac{r}{w}\right) \left[1 + \left(1 - \frac{1}{r}\right) \left(\beta^2 \frac{\mathbb{V}ar[G]}{\mathbb{V}ar[J]} - 2\beta\rho\sqrt{\frac{\mathbb{V}ar[G]}{\mathbb{V}ar[J]}}\right)\right] \quad (24)$$

In this work a latin hypercube sampling (LHS) was employed to select the initial pilot sample population. To simplify the implementation of the multilevel multifidelity variance estimate we use the same number of pilot samples to estimate both multilevel and multifidelity parameters¹. Multiple initial sample sizes are chosen and we measure the difference in the mean and standard deviation for some fixed MSE tolerance ϵ . The threshold sample size above which the difference in mean and standard deviation are within the error tolerance (ϵ) is chosen and used in subsequent calculations. Therefore the sampling is an iterative process, where we estimate the optimal number of samples at each level based on cost and observed variance. At each iteration we evaluate additional samples, which results in a more accurate variance estimation across level and model fidelities. This improved estimate is used to update the sample allocation for the next iteration. The iteration is terminated when either zero additional samples are allocated or upon reaching a prescribed maximum iteration. Note that a similar iterative sampling process is adopted for the SMLMC results presented in this work.

3 | ADJOINT SURROGATE MODEL

Motivated by the adjoint error correction of Giles and Pierce³³, Ghate² introduced an adjoint-based approximate output evaluation called the Inexpensive Monte Carlo method (IMC). The adjoint error correction is used to improve an approximation for the QoI $J(u_\delta^*, \alpha_\delta)$, where, u_δ^* is an approximation for the state u_δ due to perturbations in the design parameter α_δ as shown in Eq. (25). The linearised state equation shown in Eq. (26) is substituted in Eq. (25) to simplify the adjoint correction as shown in Eq. (27).

$$J(u_\delta, \alpha_\delta) \approx J(u_\delta^*, \alpha_\delta) - v_\delta^T \left. \frac{\partial R}{\partial u} \right|_{u_\delta^*} (u_\delta^* - u_\delta) + \mathcal{O}(\|u_\delta^* - u_\delta\|^2) \quad (25)$$

$$R(u_\delta, \alpha_\delta) \approx R(u_\delta^*, \alpha_\delta) - \left. \frac{\partial R}{\partial u} \right|_{u_\delta^*} (u_\delta^* - u_\delta) + \mathcal{O}(\|u_\delta^* - u_\delta\|^2) = 0 \quad (26)$$

$$J(u_\delta, \alpha_\delta) \approx J(u_\delta^*, \alpha_\delta) - v_\delta^T R(u_\delta^*, \alpha_\delta) + \mathcal{O}(\|u_\delta^* - u_\delta\|^2) \quad (27)$$

The adjoint v_δ^T in Eq. (27) is evaluated at the perturbed state u_δ . But one can replace the adjoint with an approximate one v_δ^{*T} and the leading error term includes the error in the adjoint solution approximation as shown below:

$$J(u_\delta, \alpha_\delta) \approx J(u_\delta^*, \alpha_\delta) - v_\delta^{*T} R(u_\delta^*, \alpha_\delta) + \mathcal{O}(\max(\|u_\delta^* - u_\delta\|^2, \|v_\delta^* - v_\delta\|^2)) \quad (28)$$

Although the IMC formulation in Eq. (28) is equivalent to the moment method² a few major differences exist. Using IMC one can estimate the complete *pdf* of the QoI J but moment methods only provide an estimate of the statistical moments like mean, variance, etc. IMC requires the evaluation of the non-linear residual R at the perturbed state, which captures non-linearities in the QoI more effectively than the moment methods. Overall in IMC there are three possible combinations to the approximation of u_δ^* and v_δ^{*T} , which leads to the IMC 1–3 formulations. The computational cost of IMC can be broken down into two parts, (i) the initial setup cost (ISC) and (ii) the evaluation cost (EC). ISC is incurred only to setup the model and is a one time cost but EC is the cost associated with every sample evaluation using IMC. In all IMC variants the EC cost is the same, which involves the

cost of calculating $J(u_\delta^*, \alpha_\delta)$ and $R(u_\delta^*, \alpha_\delta)$. Usually the cost of evaluating the objective function J is quite inexpensive compared to the evaluation of the non-linear residual R . A typical non-linear RANS solution for one sample point requires $\mathcal{O}(10^3 - 10^4)$ residual evaluations. Therefore the cost of one residual evaluation is negligible in comparison.

IMC 1: Approximates the flow and adjoint values using the unperturbed baseline solution. This method has an overall leading error of second order and it is equivalent to the first order moment method. The ISC of this method is a single adjoint solution at the mean state. Note that cost of the linear adjoint solution is independent of the number of input uncertainties and is cheaper to evaluate than the computation of the solution to the non-linear state equation, $R(u) = 0$.

$$\begin{aligned} u_\delta^* &= u \\ v_\delta^* &= v \end{aligned} \quad (29)$$

IMC 2: Approximates the flow values from the baseline tangent-linear solution and extrapolation to the perturbed state. The adjoint is approximated using the baseline solution similar to IMC 1. This method has an overall leading error of third order and it is equivalent to the second order moment method. For $\alpha \in \mathbb{R}^M$ IMC 2 requires M tangent-linear solutions in addition to the baseline non-linear and adjoint solution.

$$\begin{aligned} u_\delta^* &= u + \frac{du}{d\alpha}(\alpha_\delta - \alpha) \\ v_\delta^* &= v \end{aligned} \quad (30)$$

IMC 3: Approximates the flow and adjoint values using gradient extrapolation of the tangent-linear flow and adjoint solution. This is the most expensive, yet, most accurate method. The method has an overall leading error of order four. For $\alpha \in \mathbb{R}^M$, IMC 3 requires M tangent-linear flow and adjoint (Hessian) solutions. The tangent-linear adjoint solution is calculated using either the Tangent-on-Tangent or Tangent-on-Reverse strategy³⁴.

$$\begin{aligned} u_\delta^* &= u + \frac{du}{d\alpha}(\alpha_\delta - \alpha) \\ v_\delta^* &= v + \frac{dv}{d\alpha}(\alpha_\delta - \alpha) \end{aligned} \quad (31)$$

3.1 | Burgers' equation with uncertain boundary condition

The viscous Burgers' equation can be used as a simplified model for the displacement of shock waves in a compressible non-linear RANS model under uncertain perturbation. As a case study, consider the viscous Burgers' equation defined in the interval $x \in [-1, 1]$, with Dirichlet boundary conditions specified at the two end points $x = -1$ and $x = 1$ as shown in Eq. 32-34. An uncertain parameter δ is specified at the boundary $x = -1$, which controls the perturbation from the unit value.

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} = \nu \frac{\partial^2 u}{\partial x^2}, \quad x \in [-1, 1] \quad (32)$$

$$u(-1, t) = u(-1) = 1 + \delta \quad (33)$$

$$u(1, t) = u(1) = -1 \quad (34)$$

Xiu³⁵ gives the detailed analytic solution procedure for this problem. The steady-state solution is shown in Eq. (35), where x_0 is the location of the transition point $u(x_0) = 0$, and A is the slope $-\partial u / \partial x|_{x=x_0}$ at the transition point. This slope is obtained as a solution to Eq. 36, which is solved iteratively using a root finding algorithm. Once A is determined, x_0 is obtained by substituting A into any one of the eqs. (33)-(34). Note that the viscosity ν determines the steepness of the shock at $x = x_0$ and a fixed value of viscosity $\nu = \frac{1}{20}$ was used in all simulations.

$$u(x) = -A \tanh \left[\frac{A}{2\nu} (x - x_0) \right] \quad (35)$$

$$f(A, \delta) : (1 + \delta + A^2) \tanh \left(\frac{A}{\nu} \right) - (2 + \delta)A = 0 \quad (36)$$

The steady state solution of the viscous Burgers' equation is plotted in Fig. 3 for various values of δ . We consider the location of the transition point x_0 at steady-state (denoted by the dark circle in Fig. 3) as the QoI in this work. The value of x_0 is readily




FIGURE 3 Steady-state solution of viscous Burger's equation for indicated input perturbation δ and transition point x_0 (denoted by a filled circle) and $\nu = \frac{1}{20}$

obtained by substituting the value of A in Eq. 37.

$$x_0 = 1 - \frac{2\nu}{A} \tanh^{-1} \left(\frac{1}{A} \right) \quad (37)$$

For small values of δ the following asymptotic solution for x_0 at steady-state is shown to be valid³⁵.

$$x_0 = 1 + \nu \ln \left(\frac{\delta}{2} \right) \quad (38)$$

Note that the QoI x_0 is hypersensitive to boundary perturbation because even small changes in δ bring about huge variations in x_0 . We assume the input uncertainty δ to occur equally likely in the given interval $[a, b]$. Therefore the occurrence of δ is modelled as a random variable with a uniform probability distribution $p(\delta)$, with mean $\mu_\delta = 2^{-1}(b-a)$ and variance $\text{Var}[\delta] = 12^{-1}(b-a)^2$, shown in Eq. (39).

$$p(\delta) = \begin{cases} (b-a)^{-1} & \text{for } a \leq \delta \leq b \\ 0 & \text{otherwise} \end{cases} \quad (39)$$

For the UQ study in this work the values of $[a, b] \equiv [0, 0.1]$ is chosen. This model problem is simple, but sufficiently complex to assess the accuracy and reliability of the surrogate models.

3.2 | Approximating model response using adjoint correction

Consider the QoI J to be the transition point location x_0 and one can form the non-linear state equation $R(A, \delta)$ from Eq. (36) as,

$$R(A, \delta) \equiv (2 + \delta)A - (1 + \delta + A^2) \tanh \left(\frac{A}{\nu} \right) = 0. \quad (40)$$

For the IMC estimator, one first linearises about the mean of the input disturbance $\delta \in [0, 0.1]$, which is $\mu_\delta = 0.05$. The perturbation about this mean can be evaluated from the adjoint and tangent-linear solution using any of the IMC schemes described in Eq. (29)-(31). In Fig. 4 the QoI obtained from the various IMC estimators are plotted for a range of perturbations δ

Model	$\mathbb{E}[x_0]$ (% error)	$\sigma(x_0)$ (% error)	$\rho(x_0)$
Exact	0.80728	5.28691×10^{-2}	-
IMC 1	0.82335 (+2%)	3.17072×10^{-2} (-40%)	0.8760
IMC 2	0.78691 (-2.5%)	7.11764×10^{-2} (+35%)	0.9446
IMC 3	0.71554 (-11%)	1.71379×10^{-1} (+225%)	0.9484

TABLE 1 Comparison of mean and standard deviation between exact solution and IMC 1 – 3 methods and Pearson correlation ρ between the exact and model solution




FIGURE 4 Comparison of transition point location x_0 between exact solution and approximate IMC 1 – 3 models for a range of perturbation values $\delta \in (0, 0.1)$; zoomed view near the mean value shown on the right.

along with the exact value. The mean and standard deviation for the exact and IMC 1 – 3 models and the Pearson correlation ρ between the exact and IMC 1 – 3 are shown in Tbl. 1. The higher order adjoint surrogates IMC 2 and 3 suffer from large errors in the mean and variance compared to lower order model IMC 1 in stark contrast to what one would expect. But they predict the right trends in the solution which is reflected in the better correlation.

The results of the Burgers' model equation provide two important insights. For highly sensitive non-linear problems adjoint surrogates do not provide reliable estimates of the statistical moments (especially the variance). In fact, higher order corrections are actually detrimental. However, they do capture the mean and general trends in the function variations (better correlation). An attractive feature of the lower order adjoint surrogate IMC 1 is that the computational cost is independent of the number of input variables and requires one adjoint solution for each QoI. Higher order IMC 2/3 incur additional computational cost, which scales linearly with the number of input uncertainties for IMC 2 and quadratically for IMC 3. Therefore by using higher order IMC 2/3 we loose the attractive convergence property of classical MC; independence of convergence to the number of input uncertainties. Good correlation and low computational cost makes IMC 1 an ideal candidate for a low-fidelity model in a MFMC or MFML framework which we show in the next section.

4 | STOCHASTIC MODEL FOR MANUFACTURING UNCERTAINTIES

In this work a synthetic surface perturbation model based on GP similar to the approach presented in refs.^{36,37,3,4} has been adopted due to a lack of access to empirical blade measurements. In addition, parametrisation of the perturbations is avoided since this would lead to a reduction in the space of possible perturbed geometries³⁸. Therefore, a free node parametrisation³⁸ is used where every surface node is a random parameter defining the blade shape. The infinite dimensional probability space is approximated using a finite number of random variables, where every surface mesh node is a parameter³⁷.

Let $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$, where each $x_i \in \mathbb{R}^3$, be the set of coordinates defining the nominal surface of the blade and let \hat{n} denote the normal vector of this nominal surface. A zero-mean Gaussian process $\delta(\mathbf{x})$ is imposed along the normal direction to this nominal surface to generate the perturbed surface x_δ as,

$$\mathbf{x}_\delta = \mathbf{x} + \delta(\mathbf{x})\hat{n}. \quad (41)$$

In addition, the random variables $\delta_i = \delta(x_i)$ and $\delta_j = \delta(x_j)$ at any two arbitrary points i and j on the surface (see Fig. 5) are assumed to have a squared exponential spatial covariance as shown in Eq. (42). The squared exponential covariance has been extensively used in robust optimisation studies involving surface perturbations (see ref.^{6,39,29}). This covariance function is infinitely differentiable, which means that the GP with this covariance function has derivatives of all orders (in the mean squared sense), and is thus very smooth⁴⁰. In Eq. (42), the parameter b controls the height of the perturbations and l is the characteristic correlation length of the perturbation. A large value of l results in a larger spread of the disturbance around a given surface node. In this work, the perturbation height b and the correlation length l are assumed be a fraction of some characteristic dimension. The parameters b and l are illustrated for the surface perturbation over the LS89 cascade surface in Fig. 5. The model is used to




FIGURE 5 Surface node distribution on LS89 cascade surface along with the two arbitrary surface nodes i and j used to illustrate the spatial covariance function in Eq. 42

represent the surface variations over the LS89 turbine cascade, which is presented in the results section.

$$C_{ij} = b^2 \exp\left(-\frac{\|x_i - x_j\|^2}{2l^2}\right) \quad (42)$$

4.1 | Dimensionality reduction and goal-based truncation

The convergence properties of Monte Carlo are valid under the assumption of independent identically distributed (*i.i.d*) input parameter space³⁰. Therefore, statistical independence is an important property to have for the input variations. For a joint Gaussian distribution to be independent its covariance function must be uncorrelated i.e., the covariance matrix C_{ij} in Eq. (42) should be transformed to a diagonal matrix Λ to make the input uncertainties independent.

The expansion using the eigenvalues and vectors of the discrete zero-mean surface perturbation is shown below

$$\delta(x_i, \omega) = \sum_{i=1}^n \sqrt{\lambda_i} \mathbf{z}_i Y_i(\omega) \approx \sum_{i=1}^m \sqrt{\lambda_i} \mathbf{z}_i Y_i(\omega) \quad (43)$$

where, $Y_i(\omega)$'s are uncorrelated Gaussian random variables with zero mean and unit variance⁴¹. The expansion is usually truncated to the first m dominant eigenmodes for reducing the input dimension. Using Mercer's theorem one can show that the truncated expansion is a suitable approximation, if the eigenvalues decay sufficiently fast and m is sufficiently large³⁸. For a Gaussian covariance function the eigenvalues will exponentially decay towards zero⁴². Therefore, with a fraction of the total number of modes the variations can be approximated reliably.

The modal fraction $\bar{\lambda}_i$ and the partial modal fraction $\bar{\Lambda}_i$ shown in Eq. (44) are useful quantities to aid the truncation. Partial modal fraction can be used to estimate the value of m given a user specified threshold error. For example, if one truncates after the i^{th} PCA mode and $\bar{\Lambda}_i = 0.9$ then 90% of the total input variations (spectral content) are captured using the truncated set. A value of $\bar{\lambda}_i = 0.1$ denotes that the i^{th} PCA mode accounts for 10% of the total input variations (spectral content).

$$\bar{\lambda}_i = \frac{\lambda_i}{\sum_{j=1}^n \lambda_j}, \quad \bar{\Lambda}_i = \frac{\sum_{j=1}^i \lambda_j}{\sum_{j=1}^n \lambda_j} \quad (44)$$

The impact of an input variation on the QoI is encoded in the adjoint sensitivity. Therefore the dot product of PCA modes and the sensitivity gives the influence of the mode on the QoI. We call this scalar the mode effectiveness (η_{PCA}).

$$\eta_{PCA_i} = \lambda_i \frac{dJ}{d\mathbf{z}_i} \approx \lambda_i \frac{dJ}{dx} \cdot \mathbf{z}_i \quad (45)$$

Since the adjoint sensitivity is independent of the number of input parameters, it is evaluated once and used for all the PCA modes. The modes are sorted based on the η_{PCA} value and truncated using a new modal fraction defined in Eq. (46). Note that Eq. (46) is obtained by replacing λ with the η_{PCA} in Eq. (44). The projected PCA modes are called the goal-based PCA modes

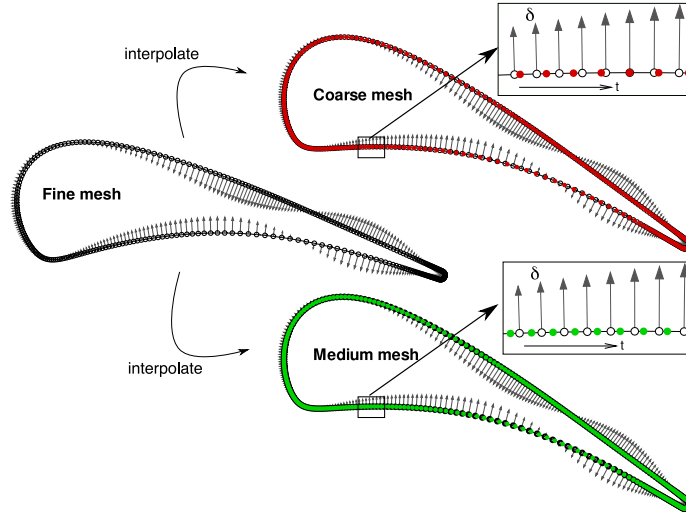


FIGURE 6 Interpolation of PCA mode #4 from the fine to the medium and coarse meshes

or simply G-PCA.

$$\bar{\lambda}_i = \frac{\eta_{PCA_i}}{\sum_{j=1}^n \eta_{PCA_j}}, \quad \bar{\Lambda}_i = \frac{\sum_{j=1}^i \eta_{PCA_j}}{\sum_{j=1}^n \eta_{PCA_j}} \quad (46)$$

Most references on G-PCA applied to surface perturbation models consider only a single QoI or a single flow condition for the truncation, for example ref. ³⁶. In this work, multiple QoIs at multiple flow conditions are considered for the G-PCA truncation (see Sec. 5.3).

4.2 | FastUQ numerical implementation

To the authors knowledge this is the first work to use MLMF method for the UQ of surface variations. Therefore, a practical implementation of this approach is provided for reference. The method is shown schematically in Fig. 7. To begin with the G-PCA modes and amplitudes are extracted from the GP model of the surface variations. The modes are typically obtained for the finest mesh (in the sequence of meshes) to maximise the space of possible perturbed geometries. The selected modes are interpolated from the surface mesh of the finest to all coarse levels. A piecewise-linear interpolation along the cascade surface (see Fig. 6) is employed in this work to interpolate the PCA modes. The local parametric coordinate t along the cascade surface used for the piecewise-linear interpolation is shown in Fig. 6. The interpolation ensures that the samples across levels belong to the same random path¹⁸.

Independent Gaussian random samples with zero-mean and unit variance are generated for the selected G-PCA modes. One can reconstruct the surface disturbance sample using the weighted summation in Eq. (43). A parallel distance-weighted interpolation (IDW) mesh smoothing is used to propagate the surface perturbations into the volume.

The perturbed mesh is then used to run either the HF or LF model to obtain the output QoI. Therefore, the total computational cost of a sample evaluation is a sum of the computational cost of mesh smoothing and QoI evaluation (using HF or LF model).

The FastUQ sampler runs the initial sampling for the LF and HF model at each level to construct the model correlations and multilevel parameters. Then based on the MSE tolerance the number of additional LF and HF samples are estimated at each level using Eq. (20). Iteratively samples are added based on the MSE estimates derived in Sec. 2.4.

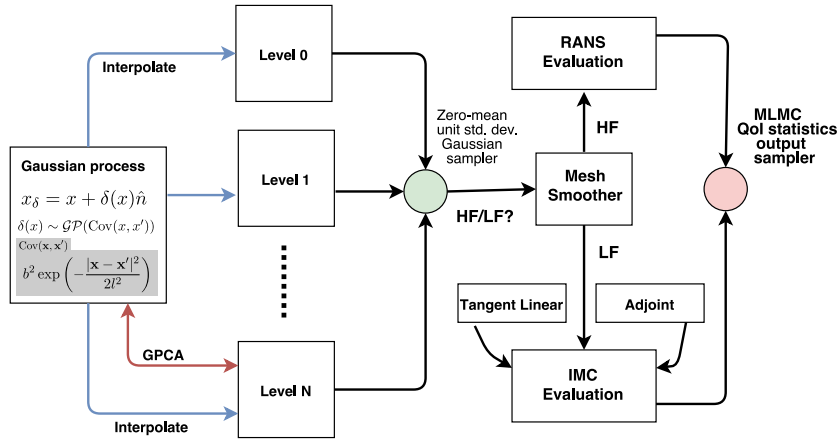


FIGURE 7 FastUQ workflow for surface variations

Quantity	Value
Pitch	57.500 mm
Axial chord	36.985 mm
Leading-edge radius	4.127 mm
Trailing-edge radius	0.710 mm
Perturbation height (b)	0.1 mm
Perturbation width (l)	10.0 mm

TABLE 2 VKI LS89 cascade geometry

Quantity	MUR43	MUR47
Inlet total-pressure (P_{01})	1.435 bar	1.596 bar
Inlet total-temperature (T_{01})	420 K	420 K
Outlet back-pressure	0.904 bar	8.236 bar
Outlet isentropic Mach (M_{isen})	0.8	1.2

TABLE 3 MUR43 and MUR47 test conditions

5 | NUMERICAL RESULTS

5.1 | Test case description

The LS89 turbine cascade was originally designed and optimised at the Von Karman Institute for Fluid Dynamics (VKI) for a subsonic isentropic exit Mach number of 0.9 using an iterative inverse method that modifies a given blade profile such that it fits an input surface velocity distribution⁴³. Experimental measurements consisting of surface pressure and wall heat transfer have been conducted for a range of inlet total-pressure and exit back pressure ratios^{44,45}.

Two test conditions, namely the MUR47 (transonic flow) and MUR43 (subsonic flow) are considered in this work. Each condition corresponds to a specific back pressure at the outlet and total temperature at the inlet. The details of the test conditions are tabulated in Tbl. 3 and Tbl. 2. The surface isentropic Mach number M_{isen} is obtained using the relation shown below

$$M_{isen} = \sqrt{\frac{2}{\gamma - 1} \left[\left(\frac{P_{01}}{P_s} \right)^{\frac{\gamma - 1}{\gamma}} - 1 \right]} \quad (47)$$




FIGURE 8 LS89 computational domain (every four mesh lines plotted for clarity)

where, P_{01} is the total-pressure at the inlet and P_s is the static pressure at a measurement station. A constant total temperature of 420 K was maintained at the inlet for all test conditions. The aerodynamic uncertainty quantification due to surface variation using FastUQ is demonstrated using the LS89 turbine cascade. The aim here is to estimate the statistics (mean and variance) of the QoIs mass flow rate J_{mass} and total pressure loss J_{ploss} defined below

$$J_{ploss} = \frac{\bar{p}_0^{in} - \bar{p}_0^{out}}{\bar{p}_0^{out} - \bar{p}^{out}}, \quad J_{mass} = \int \rho^{out}(\mathbf{u}^{out} \cdot dS) \quad (48)$$

subject to the surface perturbations on the cascade surface using the GP model described in the previous section. Note that the quantities with $(\bar{\cdot})$ are mass averaged values, p , ρ , \mathbf{u} are the fluid pressure, density, and velocity respectively and the superscript *in/out* denotes the values at the inlet/outlet of the domain.

The surface disturbance model has a height b approximately 10% of the trailing edge radius and a length l approximately equal to the leading edge diameter of the cascade. The various geometric parameters of the cascade are shown in Tbl. 2. Two levels of coarsening of the original mesh were performed maintaining a near wall spacing $y^+ \approx 1 - 2$. The fine, medium and coarse meshes have approximately 90k, 60k, and 40k nodes respectively.

5.2 | Flow and sensitivity validation

The mesh and the computational domain used for the validation is shown in Fig. 8. The surface isentropic Mach number (M_{isen}) obtained using CFD is plotted against the experimental results of Arts⁴⁴ for the MUR43 to MUR48 test conditions test conditions in Fig. 9. Good overall agreement is obtained between the experimental results and the RANS simulation.

The surface sensitivity verification is conducted using the LS89 turbine cascade test case. The MUR47 test condition is considered and the gradients of the QoI, namely exit mass flow rate and total-pressure loss are verified in this study. The gradient results obtained using adjoint, tangent-linear and finite-difference are compared using the directional derivative (Gâteaux), which needs the definition a direction vector \mathbf{h} . In this verification exercise, the surface modes obtained using the principal component analysis shown in Sec. 5.3 are considered as the surface direction vector $\mathbf{h} \equiv \mathbf{z}_j$. Then using Taylor expansion one obtains the finite-difference form (central difference) as shown below,

$$\frac{dJ}{d\mathbf{z}_j} = \frac{J(\mathbf{x} + \frac{\epsilon}{2}\mathbf{z}_j) - J(\mathbf{x} - \frac{\epsilon}{2}\mathbf{z}_j)}{\epsilon} + \mathcal{O}(\epsilon^2) \quad (49)$$

The term $J(\mathbf{x} \pm \frac{\epsilon}{2}\mathbf{z}_j)$ is obtained from the flow solution at the perturbed mesh due to the surface deformation $\mathbf{x} \pm \frac{\epsilon}{2}\mathbf{z}_j$. Choosing a small value for ϵ yields large round-off errors and using bigger values of ϵ results in larger truncation error. Thus the optimal value of ϵ is a trade-off between the round-off and truncation error. This optimal value is obtained by plotting the error between the actual gradient and the one obtained from finite difference for various step sizes. The relative error in the directional gradient (between adjoint and finite-difference) is shown for PCA modes 2,4 and 8 respectively for the total-pressure loss and exit mass flow rate cost function in Fig. 11. Comparison of the directional derivatives obtained from the optimal finite-difference step size, tangent-linear and adjoint method are plotted in Fig. 12. Good agreement is found between all three methods for the directional gradient values indicating a correct implementation.

5.3 | Goal-based truncation results

The eigenvalues of the first ten eigenmodes of the PCA and corresponding G-PCA eigenvalue of the surface perturbation over the LS89 cascade are plotted in Fig. 13 for both test conditions MUR43 and MUR47. The ranking of the first ten dominant

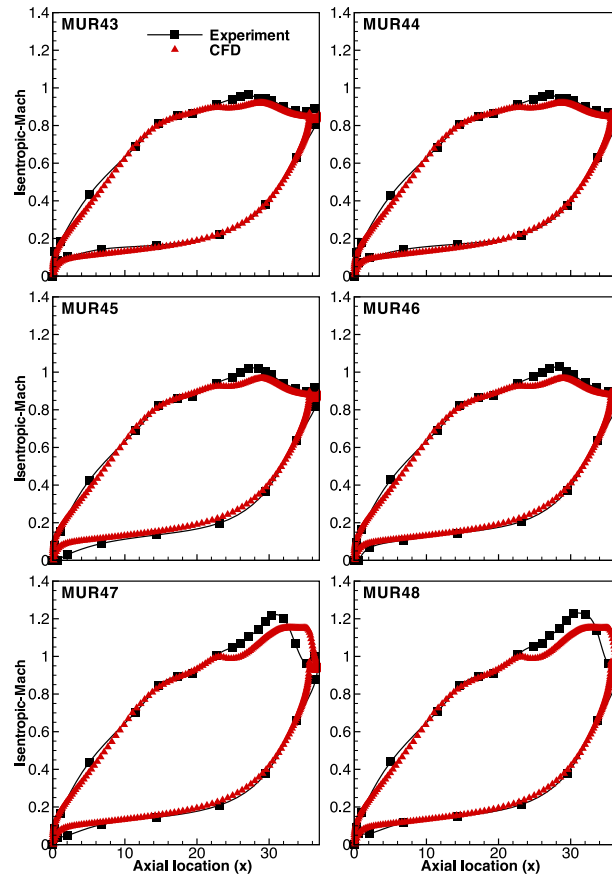


FIGURE 9 Surface distribution of isentropic Mach number for the VKI LS89 turbine cascade at indicated test conditions from ref.⁴⁴




FIGURE 10 (a) Surface displacement of PCA mode 8 and (b) generated volumetric displacement after IDW mesh smoothing (for directional derivative verification)

mode numbers are tabulated in Tbl. 4 for the exit mass flow G-PCA (EM-GPCA), and total-pressure loss G-PCA (TP-GPCA) eigenmodes. The ranking of the modes by G-PCA is quite different from the PCA. In particular, the difference is quite dramatic for the MUR47 condition. Here PCA modes number 1 and 2 are in fact the least dominant in comparison to the TP-GPCA and the mode 3 is the most dominant TP-GPCA mode. The reason becomes clear when one plots the PCA modes against the surface sensitivity field as shown in Fig. 14(a)-(l). PCA mode 2 in Fig. 14(k) has deformations occurring only in regions of very low sensitivity. Similarly, PCA modes 2 and 3 are amongst the least dominant EM-GPCA and mode number 4 is the most dominant EM-GPCA mode. In particular, PCA mode 8 is highly dominant (considering EM-GPCA) for both test conditions.




FIGURE 11 Relative error comparison of directional derivative (adjoint and finite-difference) of total-pressure loss and mass flow cost functions for modes 2, 4, and 8 (for varying step sizes).

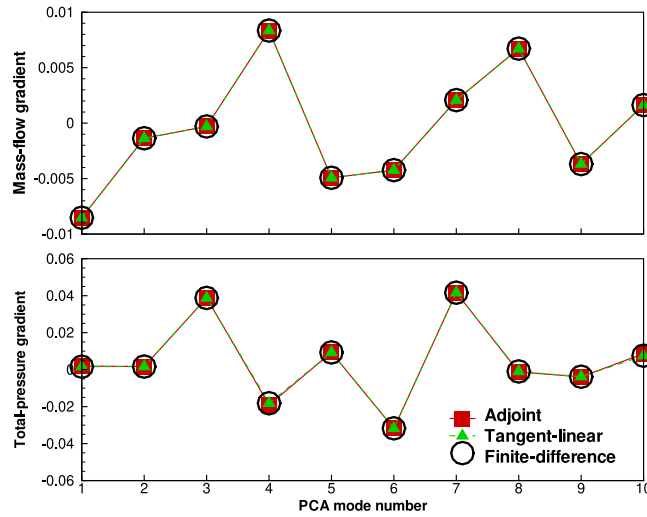


FIGURE 12 Comparison of directional derivative (adjoint, tangent-linear and finite-difference) of total-pressure loss and mass flow cost functions for PCA modes 1 to 10.

TABLE 4 Ranking of EM-GPCA and TP-GPCA modes w.r.t PCA for MUR43 and MUR47 test conditions: most-dominant (left) to least-dominant (right)

Condition	Type	PCA mode #									
		1	4	8	6	5	3	9	2	7	10
MUR43	EM-GPCA	1	4	8	6	5	3	9	2	7	10
	$\bar{\lambda}_i(\%)$	17	30	42	52	62	70	77	80	83	86
	TP-GPCA	1	4	8	5	6	9	2	7	10	13
	$\bar{\lambda}_i(\%)$	35	45	53	60	68	73	78	83	87	89
MUR47	EM-GPCA	1	4	3	7	5	9	8	11	13	2
	$\bar{\lambda}_i(\%)$	17	33	47	57	65	73	76	82	85	88
	TP-GPCA	3	7	6	4	5	11	10	15	14	1
	$\bar{\lambda}_i(\%)$	17	33	46	54	61	68	73	77	81	84

From Fig. 14(a,d) one finds that mode 8 has a constructive superposition with the sensitivity field, which makes it quite effective. Combining the (most dominant) five G-PCA modes in each QoI for the MUR43 test case i.e., {1, 4, 5, 6, 8} one can capture 50% of the variations (in the G-PCA sense). Considering the 7 modes {1, 3, 4, 5, 6, 7, 8} one can capture 50% of the spectral content of the input variations (G-PCA) considering both test conditions and QoIs. Clearly modes that are effective for a particular QoI

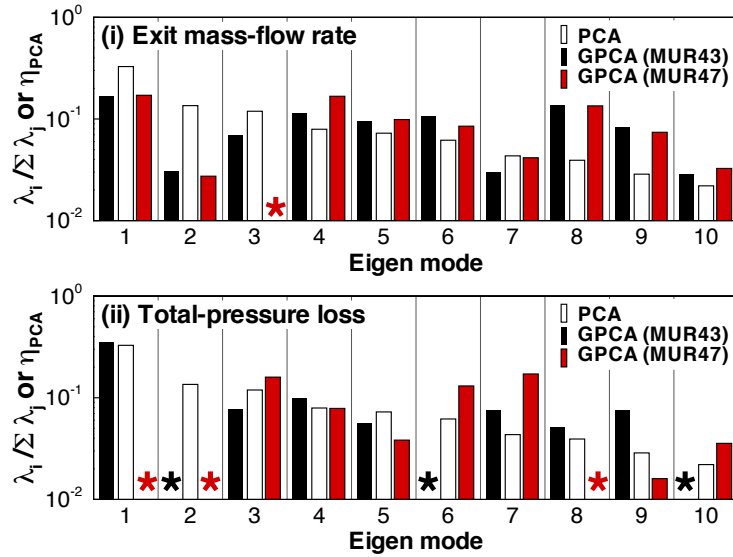


FIGURE 13 First ten PCA eigenmodes (along with the G-PCA counter-part) of the surface perturbation field over the LS89 cascade for indicated test condition and cost-function; modes with very low effectiveness ($< 10^{-2}$) are indicated by (*).




FIGURE 14 PCA eigenmode (coloured) and sensitivity (black) vector for the QoI : *exit mass flow* (red) and *total-pressure loss* (green) plotted for test conditions MUR43 (top) (a)-(f) and MUR47 (bottom) (g)-(l) for indicated mode (trailing edge sensitivity clipped for better visualisation)

and test condition need not be effective for another. Therefore, when more QoIs and test conditions are considered the number of G-PCA modes necessary to capture the spectral content of the input variations increases.

5.4 | MUR43 and MUR47 Flow Condition

The MUR43 and MUR47 flow conditions are chosen for quantifying the aerodynamic uncertainty due to surface variations. The Mach number contours for the two flow conditions are shown in Fig. 15. MUR47 has a supersonic exit isentropic Mach number of 1.02 and MUR43 has a subsonic exit isentropic Mach number of 0.84. One can notice the presence of a shock wave in MUR47 near the trailing edge of the cascade, which extends all the way into the exit. The total-pressure loss QoI is especially sensitive due to the presence of this shock. The mean and standard deviation (obtained from 20 random pilot samples) in the surface isentropic Mach number is plotted in Fig. 16. This limited sample estimate reveals the high sensitivity of the surface isentropic Mach number to the surface perturbation G-PCA modes after mid-chord of the cascade geometry on the suction side. The variations are quite high for MUR47 due to the presence of the shock aft of the cascade.




FIGURE 15 Mach number contours for the LS89 cascade at indicated test condition (for the unperturbed mesh)

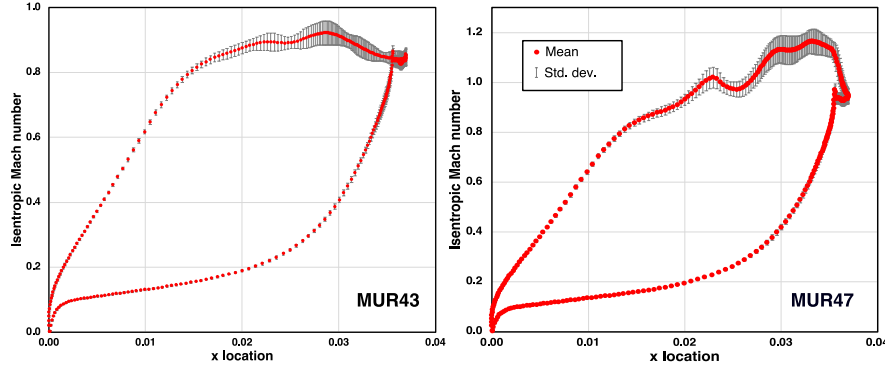


FIGURE 16 Mean and standard deviation of surface isentropic Mach number distribution at indicated flow condition ($\pm\sigma$ shown using error bars) obtained using 20 pilot samples from the surface GP.

		Coarse	Medium	Fine
Size		40k	60k	90k
Number of MPI ranks		32	32	32
Non-linear flow solver				
Computational Cost (minutes)	MUR43	15.0	28.6	85.1
	MUR47	31.8	60.7	171.6
Runtime ratio (w_l)	MUR43	0.18	0.34	1.00
	MUR47	0.19	0.35	1.00

TABLE 5 Computational cost comparison of non-linear solution

5.5 | Cost model parameters

The SMLMC and FastUQ adjoint-assisted MLMF require a cost model for the multilevel non-linear flow solution. The computational cost for the baseline (unperturbed) mesh with a residual convergence tolerance of 10^{-11} is tabulated in tbl 5. The runtime ratio w_l for each level is also shown in Tbl. 5. Note that the runtime ratio for both MUR43 and MUR47 are quite similar (see Tbl. 5). Hence a fixed runtime ratio of $\{0.2 : 0.35 : 1.0\}$ was used throughout. The computational cost for the baseline geometry (nominal or unperturbed cascade surface) shown in Tbl. 5 is obtained using a uniform flow initialisation. Note that the flow field for every HF sample evaluation in the UQ analysis is initialised using this nominal solution to reduce the number of iterations required for convergence. In addition, the number of iterations to convergence varies depending on the size of the perturbation. If the perturbations are large then more iterations are necessary and vice versa. To simplify the analysis, the runtime ratio w_l is assumed fixed based on the nominal cost and variations in computational cost due to difference in the perturbation size are ignored.

The cost of the adjoint solution and the relative cost ratio with respect to the HF solution at each level is shown in Tbl. 6. Note that the adjoint solution is significantly cheaper than the non-linear solution because the Jacobian matrix computation and

Adjoint (Exit mass flow)				
Computational cost (minutes)	MUR43	3.8	6.9	32.0
	MUR47	3.3	6.5	28.3
Cost ratio (wrt non-linear solution)	MUR43	0.3	0.2	0.4
	MUR47	0.1	0.1	0.2
Adjoint (Total pressure loss)				
Computational cost (minutes)	MUR43	2.9	5.5	30.6
	MUR47	2.8	5.1	29.2
Cost ratio (wrt non-linear solution)	MUR43	0.2	0.2	0.4
	MUR47	0.1	0.1	0.2

TABLE 6 Computational cost comparison of adjoint solution

	Tangent solution	
	MUR43	MUR47
Computational Cost (minutes)	31.3	43.6
Cost ratio (wrt non-linear solution)	0.37	0.25

TABLE 7 Computational cost comparison of tangent solution for mode #4 (fine mesh)

its ILU0 factorisation are performed once during the first iteration and reused for all subsequent iterations. But the non-linear primal requires re-computation of the Jacobian matrix and its ILU0 factorisation for every iteration. The computational cost of the tangent solution for G-PCA mode #4 for MUR43 and MUR47 flow condition is shown in Tbl. 7. Using the cost of the adjoint (see Tbl. 6) and tangent solution (see Tbl. 7) one can estimate the setup cost of the IMC 1/2 model using the Eq. (50) shown below.

$$\text{Setup cost of IMC 1} = \#QoI \times \text{AdjointCost}$$

$$\text{Setup cost of IMC 2} = \#QoI \times \text{AdjointCost} + \#\text{Modes} \times \text{TangentCost} \quad (50)$$

Note that the cost model shown in Eq. 50 is for a single flow condition whose cost of adjoint and tangent linear solution is AdjointCost and TangentCost. #QoI and #Modes in Eq. 50 denote the number of QoIs and G-PCA modes included in the analysis.

5.6 | Effect of Pilot Sampling

The accuracy of the variance estimate at each level depends on the size of the initial samples set considered, which are also called the pilot samples (see Sec. 2.5). One needs to ensure that the errors introduced by the pilot sampling is minimal because it can deteriorate the computational gain (see Eq. 24). To test the error in the pilot sampling, the number of pilot samples was doubled from 10 to 20 and the change in QoI for an MSE tolerance (ϵ) of 0.1 was recorded for the SMLMC and is shown in Tbl. 8. Latin hypercube sampling (LHS) was employed to select the initial pilot samples. The results show that the change in mean and standard deviation with the doubling of the pilot samples is negligible in comparison to the specified tolerance. Hence we used 10 pilot samples throughout this work and a maximum of 100 iterations was used for the MLMF iterative process described in Sec. 2.5. Both runs with 10 and 20 pilot samples converge within the prescribed tolerance and maximum number of iterations.

5.7 | Effects of G-PCA Truncation

To quantify the effect of G-PCA, truncation the UQ of the LS89 cascade for the MUR47 flow condition was carried out using the MLMC of Giles (see Sec. 2.3), which we refer to as the standard MLMC or SMLMC. Two sets of SMLMC simulations were run, (i) using 7 and (ii) 25 G-PCA modes (combining modes based on both QoIs and flow conditions). The 7 G-PCA modes

# pilot samples	Exit mass flow	Total-pressure loss
10	4.62286 ($\pm 2.64565 \times 10^{-2}$)	1.89269 ($\pm 1.48167 \times 10^{-1}$)
20	4.62267 ($\pm 2.64368 \times 10^{-2}$)	1.89274 ($\pm 1.48438 \times 10^{-1}$)

TABLE 8 QoI statistics with different pilot samples in SMLMC (MUR47 condition with MSE tolerance 0.1)

capture 50% and the 25 G-PCA modes capture 99% of the total variations in the input uncertainty. The cost model parameters from Sec. 5.5 are used in the current SMLMC simulation. For both sets 10 pilot samples were used per level to obtain the initial model correlations. An MSE tolerance of 0.1 was specified for the simulations. The mean and standard deviation obtained from the simulation are tabulated in Tbl. 9. Inclusion of more modes does not dramatically alter the mean value of the exit mass flow rate but the mean total-pressure loss is quite sensitive. Although the variation in the mean of the QoI (between 7 and 20 modes) is $< 2\%$, the standard deviation has a high variation in the QoI (around 36%). Therefore, all simulations henceforth use 25 G-PCA modes since it captures 99% of G-PCA variation and the predicted variations with 7 modes differed significantly from the 25 G-PCA modes.

# G-PCA modes	Exit mass flow	Total-pressure loss
7 modes	4.62604 ($\pm 2.07639 \times 10^{-2}$)	1.85356 ($\pm 1.07648 \times 10^{-1}$)
25 modes	4.62286 ($\pm 2.64565 \times 10^{-2}$)	1.89269 ($\pm 1.48167 \times 10^{-1}$)

TABLE 9 SMLMC mean and standard deviation of indicated cost function obtained using 7 and 25 dominant (combined) G-PCA modes for the MUR47 test case (MSE tolerance 0.1)

Inclusion of 25 G-PCA modes of the LS89 cascade increases the input uncertainty space quite dramatically. Note that the computational cost of IMC 1 surrogate model is independent of the number of PCA modes (or number of input uncertainties). So a large number of input uncertainties can be handled using FastUQ with IMC 1 LF model (see Sec. 3). In addition, the G-PCA can help truncate the statistically independent PCA modes that do not contribute to the change in QoI. Therefore is well suited to solve this problem.

5.8 | Optimal Resource Allocation and MSE Convergence (SMLMC)

The SMLMC using 25 G-PCA modes was run for four MSE tolerances, 1.0, 0.1, 0.05, and 0.01 respectively for the MUR47 flow condition. The resource allocation between the three mesh levels (coarse, medium, and fine) to achieve the aforesaid MSE tolerances is shown in Fig. 17 (a) (based on Eq. (20) for optimal resource allocation). Results in Fig. 17 (a) show that all computations required sample evaluation only on the coarsest mesh up to an MSE tolerance of 0.05 (excluding the 10 pilot samples at each level). This is primarily due to well correlated values of the QoI between levels i.e, the gross variations are captured by the coarse levels. This is an important requirement for multilevel methods⁴⁶. The mean and standard deviation values obtained for the various MSE error tolerance are tabulated in Tbl. 10 where it can be observed that for an MSE tolerance of 0.01 more equivalent HF samples are necessary to capture the finer details of the variations. The equivalent HF samples are still cheaper to evaluate because most of the samples are computed from the coarse level solution (see Fig. 17(a)). The convergence of the MSE with equivalent HF samples P is plotted in Fig. 17 (b), where the dotted trend-line indicates the $\mathcal{O}(P^{-1})$ convergence. The initial convergence is sharp for $\text{MSE} < 0.05$ but for larger $\text{MSE} > 0.05$ the convergence matches $\mathcal{O}(P^{-1})$. In all SMLMC runs, the maximum number of iterations (see Sec. 2.5) was set to 100 and every MSE tolerance converged within this limit. The SMLMC took close to 27h to reach convergence in the statistics for the MSE tolerance of 0.01. We estimate that calculations involving lower MSE tolerances (and large 3D test cases) demand greater computational resources and time. Hence postponed to a future study. But this should not affect our overall findings and conclusions.

5.9 | Accuracy of Adjoint Correction

To test the accuracy of the IMC LF model, the aerodynamic UQ of LS89 cascade was carried out at the MUR47 condition for the two QoIs, namely, exit-mass flow and total-pressure loss. The IMC 1 LF model requires only a single adjoint solution per

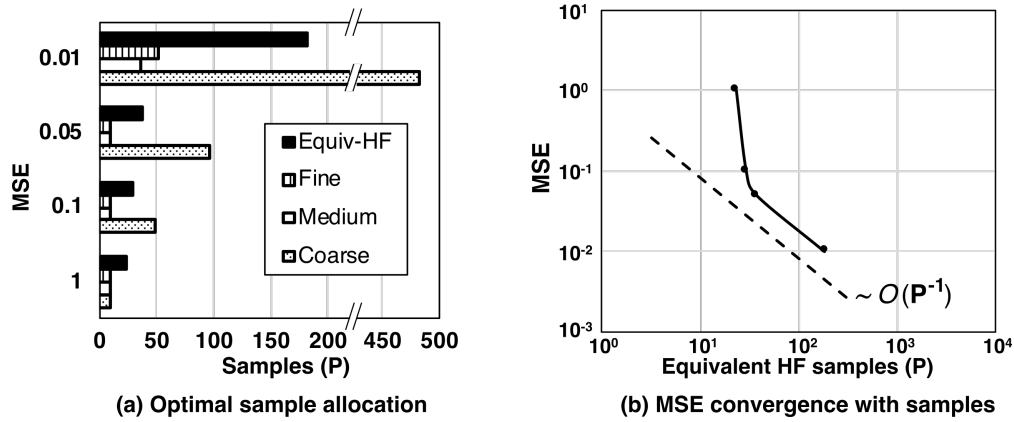


FIGURE 17 SMLMC results for the LS89 test case MUR47; (a) optimal sample allocation based on Eq. (20) and (b) MSE convergence.

MSE	Exit mass flow	Total-pressure loss	Eqv. HF
1.0	4.62043 ($\pm 3.31661 \times 10^{-2}$)	1.90907 ($\pm 1.92868 \times 10^{-1}$)	23
0.1	4.62286 ($\pm 2.64565 \times 10^{-2}$)	1.89269 ($\pm 1.48167 \times 10^{-1}$)	29
0.05	4.62339 ($\pm 2.43004 \times 10^{-2}$)	1.89003 ($\pm 1.53320 \times 10^{-1}$)	37
0.01	4.62408 ($\pm 2.42732 \times 10^{-2}$)	1.89980 ($\pm 1.47926 \times 10^{-1}$)	182

TABLE 10 SMLMC mean and standard deviation of indicated cost function obtained using 25 dominant (combined) G-PCA modes for the MUR47 test case

QoI and its computational cost is independent of the number of PCA modes (see Sec. 3). For each level the IMC 2 LF model requires as many tangent linear solutions as the number of G-PCA modes (in addition to the cost of the adjoint solution). Note that the above estimates are the number of solution evaluations to construct the LF model per level in the case of FastUQ and it is a one time cost associated with the LF model setup.

The results presented in Sec. 5.7 showed that 25 G-PCA modes were necessary to capture 99% of the input variations and using lower number of G-PCA modes gave large variations in the standard deviation of the QoI for the MUR47 test condition. The setup cost of the IMC 2 LF model for 25 modes is 17 equivalent HF samples (calculated based on Eq. (50)). This setup cost is in fact greater than the computational cost in equivalent HF samples for the entire IMC 1 simulation for an MSE tolerance of 0.01 (see Tbl. 11). Therefore, for a large number of input uncertainties the setup cost of IMC 2 becomes quite high. In addition, the MUR47 condition has a shock discontinuity in the solution. This reduces the regularity of the solution, which introduces large errors in the higher order adjoint corrections^{2,47,48} (also see Sec. 3.2) Therefore, the analysis using IMC 2/3 models is postponed to a future study and only the IMC 1 LF model is considered in all simulations.

The results of the IMC 1 are compared against the SMLMC results in Tbl. 11 for the same prescribed MSE tolerance of 0.01. Note that the sample allocation per level for SMLMC is shown in Fig. 17(a) and the IMC 1 is run on the finest mesh level. The equivalent HF samples for the IMC 1 simulation is obtained by multiplying the total runtime of the IMC 1 simulation with the ratio of the SMLMC equivalent HF samples to the SMLMC runtime (for the prescribed MSE tolerance).

Type	Exit mass flow	Total-pressure loss	Eqv. HF
SMLMC	4.62408 ($\pm 2.42732 \times 10^{-2}$)	1.89980 ($\pm 1.47926 \times 10^{-1}$)	182
IMC 1	4.62367 ($\pm 2.41368 \times 10^{-2}$)	1.77778 ($\pm 1.52438 \times 10^{-1}$)	12

TABLE 11 Mean and standard deviation of indicated cost function obtained using 25 dominant (combined) G-PCA modes for the MUR47 test case (SMLMC and IMC 1)




FIGURE 18 Pearson correlation between HF and LF model for each mesh level (MUR43 and MUR47 test condition and indicated QoI)




FIGURE 19 Percentage reduction in computational cost compared to SMLMC results.




FIGURE 20 MUR47 mean and standard deviation variation with respect to SMLMC results

Two important observations can be made from the results shown in Tbl. 11. Firstly, the IMC 1 takes only a fraction of the cost of the SMLMC ($\approx \frac{1}{10}$) for the same MSE tolerance. The IMC took 1.7 h to converge to the prescribed tolerance compared to 27 h for the SMLMC. Secondly, the total-pressure loss (mean and variance) obtained using IMC 1 has large deviations from the SMLMC results but the deviation of the exit mass flow rate QoI is lower. This clearly shows the shortcoming of the IMC 1 model for sensitive and non-linear QoIs, which can be remedied using the FastUQ method.

5.10 | FastUQ Results

The performance of the MLMF is a strong function of the correlation between the HF and IMC LF model (see Eq. (20)). Therefore, the correlation at each level for the MUR43 and MUR47 test case is plotted in Fig. 18 (obtained from the 10 pilot samples). One finds that the correlation is quite high for both MUR47 and MUR43 test conditions (average correlation for both cases is ≥ 0.95). But the coarse level in MUR43 has a lower correlation value of 0.76. Therefore, one can expect more HF sample evaluation at the coarse level for MUR43 to compensate for the lower correlation. This is clearly reflected in the resource allocation shown in Fig. 22, where 151 additional HF samples are evaluated for the MUR43.




FIGURE 21 MUR43 mean and standard deviation variation with respect to SMLMC results




FIGURE 22 Computational resource allocation across levels and models

The QoI statistics for both MUR43 and MUR47 test conditions were computed using all three UQ methods for an MSE tolerance of 0.01 and the results are summarised in Fig. 19-21. The computational cost reduction in comparison to the runtime of the SMLMC achieved by FastUQ and IMC are shown in Fig. 19. The percentage variation in the mean and standard deviation from SMLMC results for FastUQ and IMC are shown in Fig. 20 and 21. The computational cost comparison in Fig. 19 shows that FastUQ gives $\approx 70\%$ reduction in computational cost over SMLMC and the variations in mean and standard deviation of the QoI are within 1% variations of the SMLMC results. IMC achieves an additional 25% reduction compared to FastUQ. Although the mean and standard deviation for the exit mass flow rate QoI is captured within 1% tolerance, IMC has significantly large deviations for the total-pressure loss QoI. Variations as large as 10% can be observed in the standard deviation and 6% in mean values.

Based on the resource allocation shown in Fig. 22, FastUQ uses the LF model more in regions where correlations are better but resorts to running the HF model when the LF model has large deviations (lower correlation). Since the model correlations for the MUR47 test case are higher (in the range 0.98-0.99 as shown in Fig. 18), more LF samples are evaluated using the IMC 1 model and the correlation is used to correct the QoI. On the other hand, the correlation for the coarse mesh is lower in MUR43 so additional HF samples are used to compensate this error.

6 | SUMMARY

We present a framework for uncertainty quantification based on multilevel multifidelity Monte Carlo acceleration called FastUQ. Although the classical Monte Carlo does not suffer from the curse of dimensionality (CoD), multifidelity control variates can suffer from CoD due to the computational cost of constructing the low-fidelity surrogate model. In FastUQ we use a surrogate model based on adjoint sensitivity whose computational cost is independent of the number of input parameters, hence free from CoD. FastUQ is used to quantify uncertainty (mean and standard deviation) in aerodynamic parameters like total-pressure loss and mass flow rate due to manufacturing variations on a highly loaded turbine cascade. The surface variations due to manufacturing were modelled as a gaussian process and goal-based principal component analysis (PCA) considering multiple adjoint sensitives of QoIs was used for dimensionality reduction. Even with dimensionality reduction, the resulting parameter space was quite high; 25 PCA modes for a 2D cascade. Results from Lange et al.⁸ also corroborates with our findings. The mean and standard deviation of QoIs considering the PCA modes obtained using FastUQ was compared to a simple multilevel method and running a classical MC using the IMC surrogate model. FastUQ provides estimates of mean and standard deviation

with the same accuracy of MLMC and achieves 70% reduction in computational cost. For benign QoIs like mass flow rate IMC gives acceptable accuracy of mean and standard deviation with much lower computational cost but suffers from large errors for highly sensitive QoIs namely the total-pressure loss especially for non-linear flow conditions (flow with shocks). FastUQ can compensate for errors in the surrogate model by running high fidelity samples allocated optimally across multiple levels to achieve both accuracy and computational efficiency.

7 | ACKNOWLEDGEMENTS

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 642959.

8 | DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

1. Geraci G, Eldred MS, Iaccarino G. A multifidelity control variate approach for the multilevel Monte Carlo technique. tech. rep., Center for Turbulence Research, Stanford University; 2015.
2. Ghate D. *Inexpensive Uncertainty Analysis for CFD Applications*. PhD thesis. University of Oxford, 2013.
3. Hacker J. *Statistical Analysis of Manufacturing Deviations and Classification Methods for Probabilistic Aerothermal Design of Turbine Blades*. Diplomarbeit. University of Stuttgart, Stuttgart, Germany, 2000.
4. Dow EA, Wang Q. Optimal Design and Tolerancing of Compressor Blades Subject to Manufacturing Variability. In: American Institute of Aeronautics and Astronautics; 2014; Reston, Virginia: 1–11
5. Dow EA, Wang Q. The Implications of Tolerance Optimization on Compressor Blade Design. *Journal of Turbomachinery* 2015; 137(10): 7. doi: 10.1115/1.4030791
6. Dow EA, Wang Q. Simultaneous Robust Design and Tolerancing of Compressor Blades. In: ASME; 2014; Düsseldorf, Germany: 12
7. Vessaz C, Tournier C, Münch C, Avellan F. Design optimization of a 2D blade by means of milling tool path. *CIRP Journal of Manufacturing Science and Technology* 2013; 6(3): 157–166. doi: 10.1016/j.cirpj.2013.05.002
8. Lange A, Voigt M, Vogeler K, Johann E. Principal component analysis on 3D scanned compressor blades for probabilistic CFD simulation. In: American Institute of Aeronautics and Astronautics; 2012; Reston, Virginia
9. Garzon VE. *Probabilistic Aerothermal Design of Compressor Airfoils*. PhD thesis. Massachusetts Institute of Technology, 2003.
10. Worzel WP, Almal A, MacLean CD. Lifting the Curse of Dimensionality. In: Boston, MA: Springer US. 2007 (pp. 29–40)
11. Liu JS. *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics New York, NY: Springer New York . 2004
12. Xiu D, Karniadakis GE. The Wiener–Askey Polynomial Chaos for Stochastic Differential Equations. *SIAM Journal on Scientific Computing* 2002; 24(2): 619–644. doi: 10.1137/S1064827501387826
13. Adams BM, Ebeida MS, Eldred MS, et al. Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.7 User's Manual. tech. rep., Sandia National Laboratory; 2017.

14. Haldar A, Mahadevan S. *Probability, Reliability and Statistical Methods in Engineering Design*. Wiley. 1st ed. 1999.
15. Ng LWT, Willcox KE. Multifidelity approaches for optimization under uncertainty. *International Journal for Numerical Methods in Engineering* 2014; 100(10): 746–772. doi: 10.1002/nme.4761
16. Ng LWT, Willcox KE. Monte Carlo Information-Reuse Approach to Aircraft Conceptual Design Optimization Under Uncertainty. *Journal of Aircraft* 2016; 53(2): 427–438. doi: 10.2514/1.C033352
17. Peherstorfer B. Multifidelity Monte Carlo estimation with adaptive low-fidelity models. In: ; 2017.
18. Giles MB. Multilevel Monte Carlo Path Simulation. *Operations Research* 2008; 56(3): 607–617. doi: 10.1287/opre.1070.0496
19. Mishra S, Schwab C, Šukys J. Multi-level Monte Carlo Finite Volume Methods for Uncertainty Quantification in Nonlinear Systems of Balance Laws. In: Bijl H., Lucor D., Mishra S. SC., ed. *Uncertainty Quantification in Computational Fluid Dynamics. Lecture Notes in Computational Science and Engineering, vol 92* 2013 (pp. 225–294)
20. Gorodetsky AA, Geraci G, Eldred MS, Jakeman JD. A generalized approximate control variate framework for multifidelity uncertainty quantification. *Journal of Computational Physics* 2020; 408: 109257. doi: 10.1016/j.jcp.2020.109257
21. Sobol IM. On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics* 1967; 7(4): 86–112. doi: 10.1016/0041-5553(67)90144-9
22. Nobile F, Tesei F. A Multi Level Monte Carlo method with control variate for elliptic PDEs with log-normal coefficients. *Stochastic Partial Differential Equations: Analysis and Computations* 2015; 3(3): 398–444. doi: 10.1007/s40072-015-0055-9
23. Fairbanks HR, Doostan A, Ketelsen C, Iaccarino G. A low-rank control variate for multilevel Monte Carlo simulation of high-dimensional uncertain systems. *Journal of Computational Physics* 2017; 341: 121–139. doi: 10.1016/j.jcp.2017.03.060
24. Blonigan PJ, Geraci G, Rizzi F, Eldred MS. Towards an integrated and efficient framework for leveraging reduced order models for multifidelity uncertainty quantification. In: American Institute of Aeronautics and Astronautics; 2020; Reston, Virginia
25. Rumpfkeil MP, Mavriplis DJ. Efficient Hessian Calculations Using Automatic Differentiation and the Adjoint Method with Applications. *AIAA Journal* 2010; 48(10): 2406–2417. doi: 10.2514/1.J050451
26. Rumpfkeil MP. Optimizations Under Uncertainty Using Gradients, Hessians, and Surrogate Models. *AIAA Journal* 2013; 51(2): 444–451. doi: 10.2514/1.J051847
27. Engels-Putzka A, Backhaus J. Adjoint-based Shape Sensitivities for Turbomachinery Design Optimizations. In: ; 2013; Las Palmas de Gran Canaria, Spain.
28. Luo J, Liu F. Statistical evaluation of performance impact of manufacturing variability by an adjoint method. *Aerospace Science and Technology* 2018; 77: 471–484. doi: 10.1016/j.ast.2018.03.030
29. Schillings C, Schulz V. On the influence of robustness measures on shape optimization with stochastic uncertainties. *Optimization and Engineering* 2015; 16(2): 347–386. doi: 10.1007/s11081-014-9251-0
30. Robert C, Casella G. *Introducing Monte Carlo Methods with R*. New York, NY: Springer New York . 2010
31. Pasupathy R, Schmeiser BW, Taaffe MR, Wang J. Control-variate estimation using estimated control means. *IIE Transactions* 2012; 44(5): 381–385.
32. Pisaroni M, Nobile F, Leyland P. Continuation Multilevel Monte Carlo Evolutionary Algorithm for Robust Aerodynamic Shape Design. *Journal of Aircraft* 2019; 56(2): 771–786. doi: 10.2514/1.C035054
33. Pierce NA, Giles MB. Adjoint Recovery of Superconvergent Functionals from PDE Approximations. *SIAM Review* 2000; 42(2): 247–264. doi: 10.1137/S0036144598349423

34. Christakopoulos F. *Sensitivity computation and shape optimisation in aerodynamics using the adjoint methodology and Automatic Differentiation*. PhD thesis. Queen Mary University of London, 2012.
35. Xiu D, Karniadakis GE. Supersensitivity due to uncertain boundary conditions. *International Journal for Numerical Methods in Engineering* 2004; 61(12): 2114–2138. doi: 10.1002/nme.1152
36. Schillings C. *Optimal Aerodynamic Design under Uncertainties*. PhD thesis. Universität Trier, 2010.
37. Schulz V, Schillings C. Problem Formulations and Treatment of Uncertainties in Aerodynamic Design. *AIAA Journal* 2009; 47(3): 646–654. doi: 10.2514/1.37955
38. Schulz V, Schillings C. Optimal Aerodynamic Design under Uncertainty. In: Springer. 2013 (pp. 297–338)
39. Schillings C, Schmidt S, Schulz V. Efficient shape optimization for certain and uncertain aerodynamic design. *Computers & Fluids* 2011; 46(1): 78–87. doi: 10.1016/j.compfluid.2010.12.007
40. Rasmussen CE, Williams CKI. *Gaussian Processes for Machine Learning*. MIT Press . 2005.
41. Borzi A, Winckel vG. Multigrid Methods and Sparse-Grid Collocation Techniques for Parabolic Optimal Control Problems with Random Coefficients. *SIAM Journal on Scientific Computing* 2009; 31(3): 2172–2192. doi: 10.1137/070711311
42. Frauenfelder P, Schwab C, Todor RA. Finite elements for elliptic problems with stochastic coefficients. *Computer Methods in Applied Mechanics and Engineering* 2005; 194(2-5): 205–228. doi: 10.1016/j.cma.2004.04.008
43. Vandebraembussche RA, Leonard O, Nekmouche L. Subsonic and Transonic Blade Design by Means of Analysis Codes. tech. rep., AGARD, Computational Methods for Aerodynamic Design (Inverse) and Optimization; 1990.
44. Arts T, Lambert de Rouvroit M, Rutherford AW. Aero-thermal investigation of a highly loaded transonic linear turbine guide vane cascade: A test case for inviscid and viscous flow computations. tech. rep., Von Karman Institute for Fluid Dynamics; Belgium: 1990.
45. Arts T, Lambert de Rouvroit M. Aero-Thermal Performance of a Two-Dimensional Highly Loaded Transonic Turbine Nozzle Guide Vane: A Test Case for Inviscid and Viscous Flow Computations. *Journal of Turbomachinery* 1992; 114(1): 147. doi: 10.1115/1.2927978
46. Geraci G, Eldred MS, Iaccarino G. A multifidelity multilevel Monte Carlo method for uncertainty propagation in aerospace applications. In: American Institute of Aeronautics and Astronautics; 2017; Reston, Virginia
47. Giles MB. On adjoint equations for error analysis and optimal grid adaptation in CFD. In: Caughey DA, Hafez MM., eds. *Frontiers of Computational Fluid Dynamics 2006*, World scientific. 2005 (pp. 155–170)
48. Giles MB, Pierce NA. Adjoint Error Correction for Integral Outputs. In: Barth T, Deconinck H., eds. *Error Estimation and Adaptive Discretization Methods in Computational Fluid Dynamics*, Springer, Berlin, Heidelberg. 2003 (pp. 47–95)

