

RESEARCH ARTICLE

Adaptive multi-index collocation for uncertainty quantification and sensitivity analysis

John D. Jakeman¹  | Michael S. Eldred¹ | Gianluca Geraci¹ | Alex Gorodetsky²

¹Optimization and Uncertainty Quantification, Sandia National Laboratories, Albuquerque

²Department of Aerospace Engineering, University of Michigan Ann Arbor, Ann Arbor

Correspondence

John D. Jakeman, Optimization and Uncertainty Quantification, Sandia National Laboratories, Albuquerque, NM 87123.
Email: jdjakem@sandia.gov

Funding information

Defense Advanced Research Projects Agency, EQUIPS

Summary

In this paper, we present an adaptive algorithm to construct response surface approximations of high-fidelity models using a hierarchy of lower fidelity models. Our algorithm is based on multi-index stochastic collocation and automatically balances physical discretization error and response surface error to construct an approximation of model outputs. This surrogate can be used for uncertainty quantification (UQ) and sensitivity analysis (SA) at a fraction of the cost of a purely high-fidelity approach. We demonstrate the effectiveness of our algorithm on a canonical test problem from the UQ literature and a complex multiphysics model that simulates the performance of an integrated nozzle for an unmanned aerospace vehicle. We find that, when the input-output response is sufficiently smooth, our algorithm produces approximations that can be over two orders of magnitude more accurate than single fidelity approximations for a fixed computational budget.

KEYWORDS

decision making, modeling, multifidelity, sensitivity analysis, simulation, uncertainty quantification, validation

1 | INTRODUCTION

Quantifying uncertainty in simulation-based prediction of engineered systems is essential for credible certification and design. Traditional sampling-based approaches for quantifying uncertainty, such as Monte Carlo (MC), require repeated evaluations at varying parameter realizations. Consequently, MC methods are often computationally intractable for high-fidelity simulation models because they require a prohibitively large number of evaluations to obtain moderate accuracy in response statistics. In contrast, surrogate methods such as polynomial chaos,¹⁻⁴ Gaussian processes,^{5,6} low-rank decompositions,⁷⁻¹⁰ and sparse grid interpolation¹¹⁻¹⁴ can be used to build an approximation of the input-output response, often at a fraction of the cost of MC sampling. Once the surrogate has been constructed, various uncertainty quantification (UQ) tasks, such as sensitivity analysis, density estimation, etc., can then be performed on the approximation at negligible cost.*

Despite the improved efficiency of surrogate methods relative to MC sampling, building a surrogate can still be prohibitively expensive for high-fidelity simulation models. Fortunately, a selection of models of varying fidelity and

*Reduced order models can also be used to construct surrogates and have been applied successfully for UQ on many applications.¹⁵⁻¹⁷ These methods do not construct response surface approximations, but rather solve the governing equations on a reduced basis.

computational cost are typically available for many applications. For example, aerospace models span fluid dynamics, structural and thermal response, control systems, etc. Leveraging an ensemble of models can facilitate significant reductions in the overall computational cost of UQ, by integrating the predictions of quantities of interest (QoI) from multiple sources.

Multifidelity methods utilize an ensemble of models, enriching a small number of high-fidelity simulations with larger numbers of simulations from models of varying prediction accuracy and reduced cost, to enable greater exploration and resolution of uncertainty while maintaining deterministic prediction accuracy. The effectiveness of multifidelity approaches depends on the ability to identify and exploit relationships among models within the ensemble. The relationships among models within a model ensemble vary greatly, and most existing approaches focus on exploiting a specific type of structure for a presumed model sequence. For example, References 18-21 build approximations that leverage a hierarchy of models of increasing fidelity, with varying physics and/or numerical discretizations. Multilevel²² and multi-index²³ also leverage a set of models of increasing fidelity, with the additional assumption that a model sequence forms a convergent hierarchy. Such a hierarchy can be formed from a sequence of finite element discretizations to the solutions of a partial differential equation (PDE), for example, which converges to the exact solution of the governing equations as the finite element mesh is refined. Some attention has been given to building multi-fidelity approximations using models that do not admit a strict ordering of fidelity;^{24,25} however, literature in this area is limited.

In this paper, we focus on multi-level/multi-index surrogate methods, observing that access to a sequence of numerical model discretizations can be more practical than access to an ensemble of models of varying physics fidelity. Multilevel methods^{22,26,27} build the multi-fidelity surrogate as a linear combination of approximations of the differences between the QoI of models of increasing fidelity/discretization. Resources are allocated to each model level, controlled by a single hyper-parameter, specifying the level of discretization for example, in a manner that balances computational cost with increasing accuracy. In many applications, however, multiple hyper-parameters may control the model discretization, such as the mesh and time step sizes. In these situations, it may not be clear how to construct a one-dimensional (1D) hierarchy represented by a scalar hyper-parameter. To overcome this limitation, a generalization of multi-level collocation, referred to as multi-index stochastic collocation (MISC), was developed to deal with multivariate hierarchies with multiple refinement hyper-parameters.²³

In this paper, we will present an adaptive extension of MISC and use this algorithm to reduce the cost of quantifying uncertainty in the performance of a supersonic jet engine nozzle. The main contributions of this paper are:

- The extension of MISC to function approximation. The original MISC papers focused on estimation of moments using quadrature.
- The development of an adaptive multi-index collocation approach based upon sparse grid approximation. The seminal multi-index collocation paper used *a priori* estimates of model and parameter importance to allocate samples across the model ensemble. Our algorithm adaptively balances discretization error introduced by numerically solving the governing equations, response surface error, and parameter importance.
- Formulation of a multi-index method for variance based sensitivity analysis.
- Application of multi-index collocation to an engineering application of practical importance, specifically UQ of a nozzle of an unmanned aerospace vehicle.

Multi-level/multi-index surrogate methods are closely related to many multilevel/multifidelity sampling algorithms.^{24,25,28-32} These sampling algorithms leverage correlation between the outputs of multiple models to reduce the variance in statistical estimators of quantities such as expectation. This variance reduction can result in orders of magnitude reduction in the computational cost of quantifying uncertainty, but like traditional MC sampling, the error decreases slowly as the number of samples is increased which can still render this approach infeasible in some contexts. As a rough guide, multi-fidelity surrogate methods are best applied to models which with moderate dimensionality and smooth input-output responses and sampling based methods are better suited to models with large numbers of parameters and for QoI with low regularity.

The remainder of this paper is organized as follows. Section 2 introduces two motivating examples: an advection-diffusion model problem and an aero-thermo-structural analysis of the performance of a supersonic jet engine nozzle. Section 3 presents the numerical discretization used to solve the aforementioned models, and introduces the sparse grid interpolation method which is the foundation of our approach. The adaptive multi-index collocation method is presented in Section 4, and Section 5 describes how to use MISC for global sensitivity analysis. Finally, in section 6, we

highlight the strengths of the adaptive multi-index method (AMISC) as demonstrated on the advection-diffusion model problem and the engineering-scale model of the supersonic nozzle.

2 | PROBLEM SETUP

We seek to quantify uncertainty in a broad class of stochastic partial differential equations (PDE). Let $D \subset \mathbb{R}^\ell$, for $\ell = 1, 2$, or 3 and $\dim(D) = \ell$, be a physical domain; $T > 0$ be a real number; and $\Gamma \subseteq \mathbb{R}^d$, for $d \geq 1$ and $\dim(\Gamma) = d$, be the stochastic space. The PDE is defined as

$$\begin{cases} u_t(\mathbf{x}, t, \mathbf{Z}) = \mathcal{L}(u), \\ B(u(\mathbf{x}, t, \mathbf{Z})) = 0, \\ u(\mathbf{x}, 0, \mathbf{Z}) = u_0(\mathbf{x}, \mathbf{Z}), \end{cases} \quad (1)$$

where t is the time, $\mathbf{x} = (x_1, \dots, x_\ell)$ is the physical coordinate, \mathcal{L} is a (nonlinear) differential operator, B is the boundary condition operator, u_0 is the initial condition, and $\mathbf{Z} = (Z_1, \dots, Z_d) \in \Gamma$ are a set of random variables characterizing the random inputs to the governing equation. The solution $u \in V$ is a vector-valued stochastic quantity

$$u : \bar{D} \times [0, T] \times \Gamma \rightarrow \mathbb{R}^n, \quad (2)$$

for some suitable function space V and where $\bar{D} = D \times \partial D$ is the closure of the interior domain. As an example, $u(\mathbf{x}, t, \mathbf{z})$ can be a vector of temperatures, pressures, and velocities for a specific location in the domain \mathbf{x} , specific time t , and for a specific realization \mathbf{z} of the stochastic variable \mathbf{Z} .

Often, one is more interested in quantifying uncertainty in a particular functional of the solution, called the quantity of interest (QoI), rather than the full solution. Let

$$F[u] : V \rightarrow \mathbb{R}^q, \quad q > 0. \quad (3)$$

be such a function, where F is typically a continuous and bounded functional. Then, we are interested in estimating statistics of the function

$$f(\mathbf{z}) = F[u(\cdot, \mathbf{z})] : \Gamma \rightarrow \mathbb{R}^q. \quad (4)$$

which assigns the value of the quantity of interest $F[u]$ to each realization of the random variables \mathbf{z} . Expression (4) allows us to express the QoI as a function of the random variables without needed to explicitly state the dependence of the QoI on the PDE solution.

2.1 | Advection diffusion equation

As a stochastic model problem, we will use the following transient advection-diffusion PDE, in two spatial dimensions ($D = [0, 1]^2$), to highlight the strengths and weaknesses of our proposed approach,

$$\begin{aligned} \frac{\partial u}{\partial t}(\mathbf{x}, t, \mathbf{z}) + \nabla u(\mathbf{x}, t, \mathbf{z}) - \nabla \cdot [k(\mathbf{x}, \mathbf{z}) \nabla u(\mathbf{x}, t, \mathbf{z})] &= g(\mathbf{x}, t) \quad (\mathbf{x}, t, \mathbf{z}) \in D \times [0, 1] \times \Gamma \\ u(\mathbf{x}, t, \mathbf{z}) &= 0 \quad (\mathbf{x}, t, \mathbf{z}) \in \partial D \times [0, 1] \times \Gamma, \end{aligned} \quad (5)$$

with forcing $g(\mathbf{x}, t) = (1.5 + \cos(2\pi t)) \cos(x_1)$, and subject to the initial condition $u(\mathbf{x}, 0, \mathbf{z}) = 0$. Following,¹¹ we model the diffusivity k as a random field represented by the Karhunen-Loève (like) expansion (KLE)

$$\log(k(\mathbf{x}, \mathbf{z}) - 0.5) = 1 + z_1 \left(\frac{\sqrt{\pi L}}{2} \right)^{1/2} + \sum_{k=2}^d \lambda_k \phi(\mathbf{x}) z_k, \quad (6)$$

with

$$\lambda_k = \left(\sqrt{\pi L}\right)^{1/2} \exp\left(-\frac{(\lfloor \frac{k}{2} \rfloor \pi L)^2}{4}\right) k > 1, \quad \phi(\mathbf{x}) = \begin{cases} \sin\left(\frac{(\lfloor \frac{k}{2} \rfloor \pi x_1)}{L_p}\right) & k \text{ even} , \\ \cos\left(\frac{(\lfloor \frac{k}{2} \rfloor \pi x_1)}{L_p}\right) & k \text{ odd} , \end{cases} \quad (7)$$

where $L_p = \max(1, 2L_c)$, $L = \frac{L_c}{L_p}$ and $L_c = 0.5$.

In this example we truncate (6) at $d = 10$ terms and choose Z_k , $k \in \{1 \dots, d\}$ to be uncorrelated independent and identical uniform random variables distributed on $[-\sqrt{3}, \sqrt{3}]$. Using 10 terms captures 99.9999996 percent of the total energy of the untruncated expansion. We choose a random field which is effectively 1D so that the error in the finite element solution is more sensitive to refinement of the mesh in the x_1 -direction than to refinement in the x_2 -direction.

The advection diffusion Equation (5) is solved using linear finite elements and implicit backward-Euler timestepping. In the following we will show how solving the PDE with varying numbers of finite elements and timesteps can reduce the cost of approximating the quantity of interest

$$f(\mathbf{z}) = \int_D u(\mathbf{z}) \frac{1}{2\pi\sigma^2} \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}^*\|_2^2}{\sigma^2}\right) d\mathbf{x}, \quad (8)$$


where $\mathbf{x}^* = (0.3, 0.5)$ and $\sigma = 0.16$. The advection diffusion model, defined here, can be inexpensively evaluated and therefore allows for exhaustive exploration of the behavior of our proposed multi-index collocation approach.

2.2 | Aero-thermo-structural model of a jet engine nozzle

One of the main objectives of this paper is to quantify uncertainty in a supersonic jet engine nozzle. Realistic modeling of the nozzle requires a multiphysics approach that couples aerodynamic, structural and thermal analyses. In this paper we use the model introduced in Reference 33, and Figure 1 provides a conceptual diagram of its components. This model performs a steady-state analysis of top-of-climb flight, where temperatures and stresses are the highest. Forty random variables were identified and characterized from experimental data, simulations, or expert judgment. Each variable and its distribution is described in Appendix A.

We have developed an automated suite of coupled multidisciplinary analysis tools for the steady aero-thermal-structural analysis of a supersonic nozzles called MULTI-F. MULTI-F is written in Python and calls the open source codes SU2³⁴ and AERO-S³⁵ for fluid and thermal/structural analyses, respectively. The internal and external flow is modeled by Euler equations which are solved using SU2. One-way feed-forward coupling is used to model the interaction between the fluid and structural analyses. Specifically the temperature obtained from the aerodynamic analysis furnishes the boundary condition on the inner wall of the thermal model. The pressure and temperatures obtained from the aerodynamic and thermal analyses then are used to provide the inner wall pressure and the temperature distribution used in a linear structural analysis. An elastostatic boundary value problem representing the equilibrium of internal and external forces is considered for the structural analysis while a Poisson boundary value problem representing the steady state heat transfer is considered for the thermal analysis. The other boundary conditions are state-independent; a convection boundary condition is specified on the outer wall of the thermal model, while a fixed displacement boundary condition is imposed on the outer edges of the baffles and stringers in the structural model.

FIGURE 1 (Left) schematic of the nozzle components. (Right) Component numbering. These figures can be found in Reference 33 [Color figure can be viewed at wileyonlinelibrary.com]



| Number | Component Name |
|--------|-------------------|
| 1 | thermal layer |
| 2 | inner load layer |
| 3 | middle load layer |
| 4 | outer load layer |
| 5 | stringers |
| 6-10 | baffles 1-5 |

In the following we will leverage different resolutions of both the CFD and structural meshes to reduce the cost of quantifying uncertainty in four model QoI: mass, thrust, load-layer temperature failure ratio, and thermal-layer failure criteria.

3 | MULTIFIDELITY MODELING

In this paper, we present a method to efficiently estimate statistics of the QoI (4) obtained from models of complex physical systems (1). For most, if not all, practical applications, the solutions of the governing equations and statistics of the resulting QoI cannot be computed analytically and numerical approximations must be employed. In this section, we discuss the impact of numerical simulation errors and response surface errors on the quantification of uncertainty.

Given a set of governing Equation (1), we assume access to a PDE solver that approximates the solution u of these equations for a given *fixed* \mathbf{z} . We also assume this solver has a set of hyper-parameters—mesh size, time step, maximum number of iterations, convergence tolerance, etc.—which can be used to estimate the QoI f at varying accuracy and cost. Differing settings for these hyper-parameters produces simulations of *varying fidelities* (*resolution*). We refer to approaches that leverage only one model or solver setting as *single-fidelity* methods and approaches that leverage multiple models and settings as *multi-fidelity* methods.

As stated in the introduction, quantifying uncertainty in numerical simulations can be costly, consequently it is common to construct a surrogate (function approximation) of the model parameters-to-output map. In this section, we introduce the solver parameters used to solve the two models considered in this paper (Section 3.1) and then present the method used to build surrogates of a single fidelity model (Section 3.2). These foundations are used to construct a multi-fidelity surrogate in Section 4.

3.1 | Physical approximation

The accuracy and cost of a model is often determined by a set of solver parameters, such as mesh size, time step, tolerances of numerical solvers. Let n_α denote the number of such parameters for a given model. Furthermore let $l_{\alpha_i} \in \mathbb{N}$ denote the number of values that each solver parameter can assume, such that the multi-index $\alpha = (\alpha_1, \dots, \alpha_{n_\alpha}) \in \mathbb{N}^{n_\alpha}$, $\alpha_i \in \mathcal{M}_{\alpha_i} = \{1, \dots, l_{\alpha_i}\}$ can be used to index specific values of the solver parameters. Finally let $\hat{f}_\alpha(\mathbf{z})$ to denote a single-fidelity physical approximation of the QoI $f(\mathbf{z})$ using the solver parameters indexed by α . In this section we detail the solver indices α for the two models considered in this paper and describe their effect on the accuracy and cost of $\hat{f}_\alpha(\mathbf{z})$.

3.1.1 | Advection-diffusion model

The advection-diffusion Equation (5) is solved using the Galerkin finite element method with linear finite elements and the implicit backward-Euler time-stepping scheme. Let h_1 and h_2 denote the mesh-size in the spatial directions x_1 and x_2 . Then given some base coarse discretization $h_{j,0}$, we create a sequence of l_{α_j} uniform triangular meshes of increasing fidelity by setting $h_{\alpha_j} = h_{j,0} \cdot 2^{-\alpha_j}$, $0 \leq \alpha_j < l_{\alpha_j}$, where $h_{j,0} = \frac{1}{4}$ and $l_{\alpha_j} = 6$, $j = 1, 2$. The number of vertices in the spatial mesh is $(h_{\alpha_1}^{-1} + 1)(h_{\alpha_2}^{-1} + 1)$.

The backward-Euler scheme used to evolve the advection-diffusion equation in time is an implicit method and unconditionally stable; consequently, the time-step size Δt can be made very large. In this paper, we use an ensemble of l_{α_3} time-step sizes $\Delta t_{\alpha_3, \alpha_3} = \Delta t_{\alpha_3, 0} \cdot 2^{-\alpha_3}$, $0 \leq \alpha_3 < l_{\alpha_3}$ to solve (5) and set $l_{\alpha_3} = 6$ and $\Delta t_{\alpha_3, 0} = 1/4$. For high enough α_1 and α_2 , the ensemble represents a sequence of models of increasing accuracy.

The cost of solving the advection-diffusion is dependent on both the number of mesh vertices and the number of time-steps. The number of mesh vertices is $2^{(\alpha_1+2)(\alpha_2+2)}$, whereas the number of time-steps is 2^{α_3+2} . In the following we use a conservative estimate for the computational cost of evaluating the advection diffusion model. Specifically we assume that the cost is $O(2^{(\alpha_1+2)(\alpha_2+2)(\alpha_3+2)})$. This cost model is representative of the best scaling that one could hope for when solving the advection diffusion on a large scale. It mimics the use of a linear solver whose costs grows linearly with the number of degrees of freedom and a finite element assembly which is constant regardless of the mesh discretization.




FIGURE 2 The error in $\mathbb{E}[f]$ for varying mesh and temporal resolutions. The legend labels denote the mesh discretization parameter values $(\alpha_1, \alpha_2, \alpha_3)$ used to solve the advection diffusion equation. Numeric values represent discretization parameters that are held fixed while the symbol \cdot denotes that the corresponding parameter is varying. The dashed lines represent the theoretical rates of the convergence of the deterministic error when refining h_1 (left), h_2 (middle), and Δt (right)

Figure 2 depicts the changes induced in the error of the mean of the QoI, that is, $\mathbb{E}[f]$, as the mesh and temporal discretizations are changed.[†] The error decreases quadratically with both h_1 and h_2 and linearly with Δt until a saturation point is reached. These saturation points occur when the error induced by a coarse resolution in one mesh parameter dominates the others. For example the left plot shows that when refining h_1 the final error in $\mathbb{E}[f]$ is dictated by the error induced by using the mesh size h_2 , provided Δt is small enough. Similarly the right plot shows, that for fixed h_2 , the highest accuracy that can be obtained by refining Δt is dependent on the resolution of h_1 . In Section 6.1, we leverage this varying cost-vs-accuracy profile to efficiently estimate uncertainty in the advection-diffusion QoI.

3.1.2 | Aero-thermal-structural model for a jet engine nozzle

The nozzle model has two solver parameters which control accuracy and cost. These parameters specify the mesh resolutions for the fluid and structural analyses. The internal and external nozzle flow is calculated using SU2, an open-source software suite for multiphysics simulations.³⁴ The governing Euler equations are discretized in SU2 using a finite volume method with a standard edge-based data structure. The first discretization index $0 \leq \alpha_1 < 3$ indexes the resolution of the CFD mesh. Specifically the meshes indexed by α represent a sequence of three increasingly refined meshes. The coarsest CFD mesh is depicted in the left of Figure 3.

The nozzle thermal and structural analyses are calculated using AERO-S, an open-source finite element method (FEM) analysis software.³⁵ The second discretization index $0 \leq \alpha_2 < 5$ of the nozzle model indexes a sequence of increasingly refined structural meshes. The coarsest structural mesh is depicted on the right of Figure 3.

The relative computational cost of each nozzle model approximation is given in Table 1. The cost is given in wall time using 4 CPUs for the CFD model and 1 CPU for the structural model (the structural model cannot currently be run in parallel).[‡]

3.2 | Stochastic approximation

Given evaluations of $\hat{f}_\alpha(\mathbf{z})$ at various realizations of the random parameters \mathbf{z} , this section describes a method to construct surrogates that can interpolate these samples. With this goal let $\beta = (\beta_1, \dots, \beta_{n_\beta}) \in \mathbb{N}^{n_\beta}$ denote a multi-index that controls the computational cost of constructing a surrogate and its corresponding accuracy. For example, β_i could be the number of univariate samples in the i th variable dimension, $1 \leq i \leq n_\beta = d$, used to construct a tensor product interpolant.

[†]The expectation is computed using the same ten samples for all model resolutions. This allows the error in the statistical estimate induced by small numbers of samples to be ignored. The reference solution is obtained using the model indexed by (6, 6, 6).

[‡]The timings in Table 1 were obtained with 4 Intel(R) Xeon(R) CPU E5-2640 v4 @ 2.40GHz (3.40GHz with turbo boost) processors and 64GB of RAM. The absolute numbers are not important, only the relative timings.




FIGURE 3 (Left) The coarse ($\alpha_1 = 0$) 2D CFD mesh for the nozzle geometry and pressure contours for the fluid moving from left to right. (Right) The coarse mesh ($\alpha_2 = 0$) of the structural model with geometry specified using the three-dimensional parameterization. The fluid moves from left to right

| α_1 | α_2 | 0 | 1 | 2 | 3 | 4 |
|------------|------------|-------|-------|-------|-------|-------|
| 0 | | 36.3 | 39.3 | 44.2 | 61.5 | 131.2 |
| 1 | | 88.5 | 90.5 | 95.7 | 113.1 | 184.2 |
| 2 | | 293.6 | 297.4 | 300.0 | 317.5 | 384.5 |

TABLE 1 Computational cost (in seconds on 4 CPUs) of the nozzle model

We refer to the entries of β as stochastic approximation hyper-parameters and denote the surrogate, built with this choice of parameters, as $\hat{f}_{\alpha, \beta}$.

Numerous methods, such as those discussed in the introduction, have been developed to build surrogates of simulation models. In this paper we choose to focus on sparse grid surrogates because the single-fidelity formulation, presented in this section, can easily be extended to use multiple model fidelities as shown in Section 4.

3.2.1 | Tensor-product interpolation

In this section, we describe tensor-product interpolation. Knowledge of tensor-product interpolation is an essential component of sparse grid approximation. Let $\hat{f}_{\alpha, \beta}(\mathbf{z})$ be an M -point tensor-product interpolant of the function \hat{f}_{α} , with $n_{\beta} = d$. This interpolant is a weighted linear combination of tensor-product of univariate Lagrange polynomials

$$\phi_{i,j}(z_i) = \prod_{k=1, k \neq j}^{m_{\beta_i}} \frac{z_i - z_i^{(k)}}{z_i^{(j)} - z_i^{(k)}}, \quad i \in \{1, \dots, d\}, \quad (9)$$

defined on a set of univariate points $z_i^{(j)}, j \in \{1, \dots, m_{\beta_i}\}$.[§] Specifically the multivariate interpolant is given by

$$\hat{f}_{\alpha, \beta}(\mathbf{z}) = \sum_{\mathbf{j} \leq \beta} \hat{f}_{\alpha}(\mathbf{z}^{(\mathbf{j})}) \prod_{i=1}^d \phi_{i,j_i}(z_i). \quad (10)$$

The partial ordering $\mathbf{j} \leq \beta$ is true if all the component wise conditions are true.

Constructing the interpolant requires evaluating the function \hat{f}_{α} on the grid of points

$$\mathcal{Z}_{\beta} = \bigotimes_{i=1}^d \mathcal{Z}_{\beta_i}^i = [\mathbf{z}^{(1)} \dots \mathbf{z}^{(M_{\beta})}] \in \mathbb{R}^{d \times M_{\beta}}. \quad (11)$$

[§]For tensor-product approximation and sparse grid approximation we can use n_{β} and d interchangeably because $n_{\beta} = d$.

We denote the resulting function evaluations by $F_{\alpha,\beta} = \hat{f}_\alpha(\mathcal{Z}_\beta) = [\hat{f}_\alpha(\mathbf{z}^{(1)}) \dots \hat{f}_\alpha(\mathbf{z}^{(M_\beta)})]^T \in \mathbb{R}^{M_\beta \times q}$, where the number of points in the grid is $M_\beta = \prod_{i=1}^d m_{\beta_i}$.

It is often reasonable to assume that, for any \mathbf{z} , the cost of each simulation is constant for a given α . So letting W_α denote the cost of a single simulation, we can write the total cost of evaluating the interpolant $W_{\alpha,\beta} = W_\alpha M_\beta$. Here, we have assumed that the computational effort to compute the interpolant once data has been obtained is negligible, which is true for sufficiently expensive models \hat{f}_α . In this paper, we use the nested Clenshaw-Curtis points

$$z_i^{(j)} = \cos\left(\frac{(j-1)\pi}{m_{\beta_i}}\right), \quad j \in \{1, \dots, m_{\beta_i}\}, \tag{12}$$

to define the univariate Lagrange polynomials. The number of points $m(l)$ of this rule grows exponentially with the level l , specifically $m(0) = 1$ and $m(l) = 2^l + 1$ for $l \geq 1$. We remark here that any univariate sequence of points can be used to construct a sparse grid. However, the most efficient rules are nested and tailored to the distribution of the random variables \mathbf{Z} . Clenshaw-Curtis points are popular for uniform variables and weighted Leja sequences³⁶ can be used for any continuous bounded variable as well as for a broad class of unbounded variables.

When using surrogates for uncertainty quantification, it is often useful to compute the expectation of the approximation. The expectation $\mu_{\alpha,\beta}$ of a tensor product interpolant can be computed without explicitly forming the interpolant and is given by

$$\mu_{\alpha,\beta} = \int_\Gamma \sum_{j \leq \beta} \hat{f}_\alpha(\mathbf{z}^{(j)}) \prod_{i=1}^d \phi_{i,j_i}(z_i) \, d\omega(\mathbf{z}) = \sum_{j \leq \beta} \hat{f}_\alpha(\mathbf{z}^{(j)}) v_j. \tag{13}$$

The expectation is simply the weighted sum of the Cartesian-product of the univariate quadrature weights $v_j = \prod_{i=1}^d \int_{\Gamma_i} \phi_{i,j_i}(z_i) \, d\omega(z_i)$, which can be computed analytically.

3.2.2 | Sparse grid interpolation

The computational effort required to generate function values for a tensor-product interpolation is proportional to the number of grid points, which grows exponentially with the number of random parameters d . In this section we introduce sparse grid interpolation which can construct multivariate function approximations with accuracy comparable to that of tensor-product approximations, but at a fraction of the cost.

The sparse grid interpolant and expectation of the function \hat{f}_α can be expressed as a linear combination of low-resolution tensor-product interpolants³⁷

$$\hat{f}_{\alpha,\mathcal{I}}(\mathbf{z}) = \sum_{\beta \in \mathcal{I}} c_\beta \hat{f}_{\alpha,\beta}(\mathbf{z}) \quad \mu_{\alpha,\mathcal{I}} = \sum_{\beta \in \mathcal{I}} c_\beta \mu_{\alpha,\beta}, \tag{14}$$

where the coefficients c_β are given by

$$c_\beta = \sum_{\mathbf{j} \in \otimes_{i=1}^d \{0,1\}} (-1)^{|\mathbf{j}|_1} \chi_{\mathcal{I}}(\beta + \mathbf{j}), \quad \chi_{\mathcal{I}}(\mathbf{k}) = \begin{cases} 1 & \mathbf{k} \in \mathcal{I} \\ 0 & \text{otherwise} \end{cases}. \tag{15}$$

Recalling that \mathcal{Z}_β are the grid points needed to build the tensor product interpolation $\hat{f}_{\alpha,\beta}$, then the points in a sparse grid are $\mathcal{Z}_{\mathcal{I}} = \cup_{\beta \in \mathcal{I}} \mathcal{Z}_\beta$. When using nested univariate sequences to construct the tensor product grids we can write

$$\mathcal{Z}_{\mathcal{I}} = \cup_{\beta \in \mathcal{I}} \mathcal{Z}_\beta^{\text{diff}}, \quad \mathcal{Z}_\beta^{\text{diff}} = \mathcal{Z}_\beta \setminus \cup_{\substack{\beta^* \in \mathcal{I} \\ \beta^* \leq \beta}} \mathcal{Z}_{\beta^*}, \quad \|\beta\|_1 > 0,$$

where $\mathcal{Z}_0^{\text{diff}} = \emptyset$ and $\mathcal{Z}_\beta^{\text{diff}} \cap \mathcal{Z}_{\beta^*}^{\text{diff}} = \emptyset$ if $\beta \neq \beta^*$. This decomposition of the sparse grid points indicates that many sparse grid points are shared among its constituent tensor-product grids, and thus the total number of points in the sparse grid

| d | 2 | 5 | 10 | 20 | 40 |
|------------|---------|---------|---------|---------|---------|
| $M_{I(l)}$ | 2.5e+01 | 3.1e+03 | 9.8e+06 | 9.5e+13 | 9.1e+27 |
| M_β | 1.3e+01 | 6.1e+01 | 2.2e+02 | 8.4e+02 | 3.3e+03 |

TABLE 2 The number of points $M_{I(l)}$ in a isotropic sparse grid and the number of points M_I in a tensor-product for increasing number of random parameters and fixed $l = 2$

satisfies

$$M(\mathcal{I}) = \text{card}(\mathcal{Z}_{\mathcal{I}}) = \text{card}\left(\bigcup_{\beta \in \mathcal{I}} \mathcal{Z}_\beta\right) \leq \sum_{\beta \in \mathcal{I}} M_\beta.$$

The index set \mathcal{I} used in a sparse grid approximation (14) can be tailored to minimize error whilst also minimizing cost, that is, the number of function evaluations. The only restriction on the index set is that it must be downward closed. An index set \mathcal{I} is downward closed if $\beta_1 \in \mathcal{I}$ implies that $\beta_2 \in \mathcal{I}$ for all $\beta_2 \leq \beta_1$.

Finding an efficient index set \mathcal{I} can be cast as an optimization problem. With this goal, let

$$\Delta E_{\alpha,\beta} = \|\hat{f}_{\alpha, \mathcal{I} \cup \{\beta\}} - \hat{f}_{\alpha, \mathcal{I}}\| \quad \Delta W_{\alpha,\beta} = |\text{Work}[\hat{f}_{\alpha, \mathcal{I} \cup \{\beta\}}] - \text{Work}[\hat{f}_{\alpha, \mathcal{I}}]|, \quad (16)$$

respectively denote the norm of the difference between the sparse grid with and without the interpolant $f_{\alpha,\beta}$ and the work needed to add $f_{\alpha,\beta}$ to the sparse grid. Noting that the error in the sparse grid satisfies $\|f - \hat{f}_{\alpha, \mathcal{I}}\| \leq \sum_{\beta \notin \mathcal{I}} \Delta E_{\alpha,\beta}$, we can formulate finding a quasi-optimal index set as a binary knapsack problem³⁸

$$\max \sum_{[\beta] \in \mathbb{N}_0^{n_\beta}} \Delta E_{\alpha,\beta} \delta_\beta \quad \text{such that} \quad \sum_{[\beta] \in \mathbb{N}_0^{n_\beta}} \Delta W_{\alpha,\beta} \delta_\beta \leq W_{\max}, \quad \delta_\beta \in \{0, 1\}, \quad (17)$$

for some maximum work W_{\max} . The solution $\mathcal{I} = \{\beta | \delta_\beta = 1\}$ to this problem balances the computational work of adding a specific interpolant with the reduction in error that would be achieved.[¶]

One approach for solving the knapsack problem is to use estimates of the error $\Delta E_{\alpha,\beta}$ derived from assumptions on the smoothness of the function f_α . The resulting grid depends on the space of functions that f_α belongs to, and not on the function itself. Assuming that the function f_α has continuous mixed derivatives of order r in each random variable dimension and $\Delta W_{\alpha,\beta} = \text{card}(\mathcal{Z}_\beta^{\text{diff}})$, yields the traditional isotropic sparse grid index set and corresponding sparse grid coefficients given below

$$\mathcal{I}(l) = \{\beta | (\max(0, l-1) \leq \|\beta\|_1 \leq l-d-2) \quad c_\beta = (-1)^{l-\|\beta\|_1} \binom{d-1}{l-\|\beta\|_1}. \quad (18)$$

Note again we have used the fact $n_\beta = d$ and that the form of the isotropic index set and coefficients reflect our choice to set $\beta_i \geq 0$.[#] The sparse grid using (18) is often referred to as the Smolyak algorithm.³⁹

There is no exact formula for the number of points in an isotropic sparse grid. However, for fixed level l and large enough d , the number of points in the Clenshaw-Curtis isotropic sparse grid satisfies

$$M_{I(l)} \approx \frac{2^l}{l!} d^l,$$

where we use \approx to denote the strong equivalence of sequences.⁴⁰ This is much smaller than the number of points in the tensor-product grid

$$M_I = m_I^d = (2^l + 1)^d \quad \mathbf{l} = (l, \dots, l).$$

In Table 2, we compare the number of points $M_{I(l)}$ in a isotropic sparse grid and the number of points M_I in a tensor-product grid for increasing number of random parameters and fixed $l = 2$. It is clear that the number of points in the sparse grid increases much more slowly than the number of points in the tensor-product grid.

[¶]Note \mathcal{I} is only ‘‘quasi’’ optimal since the error decomposition maximized in the knapsack problem is only an upper bound for the true error.

[#] Much of the literature assumes $\beta_i > 0$, in which case $\mathcal{I}(l) = \{\beta | (l+1 \leq \|\beta\|_1 \leq l+d)\}$ and $c_\beta = (-1)^{l+d-\|\beta\|_1} \binom{d-1}{l+d-\|\beta\|_1}$.




FIGURE 4 (Left) the function $\hat{f}_\alpha(\mathbf{z}) = \cos(2\pi z_1) \cos(\pi z_2)$. (Middle) the tensor product interpolant $\hat{f}_{\alpha,(2,2)}(\mathbf{z})$. (Right) the isotropic sparse grid interpolant $\hat{f}_{\alpha,I(2)}(\mathbf{z})$. The black circles depict the samples used to build each interpolant

For a function with r continuous mixed derivatives, the isotropic level- l sparse-grid based on Clenshaw-Curtis abscissas with $M_{I(l)}$ points satisfies:³⁷

$$\|f - \hat{f}_{I(l)}\|_{L^\infty(\Gamma)} \leq C_{d,r} M_{I(l)}^{-r} (\log M_{I(l)})^{(r+2)(d-1)+1}, \tag{19}$$

and the tensor-product interpolant satisfies

$$\|f - \hat{f}_I\|_{L^\infty(\Gamma)} \leq K_{d,r} M_I^{-r/d}, \tag{20}$$

where the constants $C_{d,r}$ and $K_{d,r}$ depend on the smoothness r and the dimension d , but not on the number of points in the grid. From (19), we can see that the error in a sparse grid approximation depends only weakly on the dimension d . In contrast, the error in a tensor product interpolant (20) depends exponentially on the number of random parameters d . For both sparse grid and tensor product interpolation, the error in each approximation improves as the smoothness r increases.

We remark here that, although an improvement over tensor-product interpolation, the logarithmic term in the sparse grid error bound can become large when the number of model variables becomes sufficiently large. In these situations, adaptive strategies for constructing the index set I have proven to be very effective. We present one such adaptive strategy in Section 4.2.

3.2.3 | An example of isotropic sparse grid approximation

In this section we use a simple example to provide insight into the construction and utility of sparse grid approximations. Consider the function^{||}

$$\hat{f}_\alpha(\mathbf{z}) = \cos(2\pi z_1) \cos(\pi z_2), \tag{21}$$

which depends on two random variables $\mathbf{Z} = (Z_1, Z_2)$ which are uniformly distributed in $\Gamma = [-1, 1]^2$. The contours of this function are depicted in Figure 4. Figure 4 also plots the points used to build a $l = 2$ tensor-product approximation, the points used to construct a $l = 2$ sparse grid approximation, and the contours of the resulting approximations.

The isotropic sparse grid, depicted in the right plot of Figure 4, is the weighted linear combination of the tensor-product interpolants shown in Figure 5. The sparse grid points are the union of the points of the low resolution grids with nonzero coefficients c_β . The sparse grid achieves a similar accuracy to the tensor-product interpolant $\hat{f}_{\alpha,(2,2)}(\mathbf{z})$ (see the middle plot of Figure 4 and the top right plot of Figure 5) but does not need to use the unique points $\mathbf{z}_\beta^{\text{diff}}$ (depicted by green squares) in the finest tensor-product grids (indexed by (1, 2), (2, 1), and (2, 2)). Moreover despite the sparse grid being comprised of low-resolution/low-accuracy tensor-product interpolants, the linear combination of these interpolants produces a much more accurate approximation than any single low-resolution interpolant. Avoiding the construction of fine tensor-product

^{||}Note the function in Equation (21) has no physical solver parameters, however, we keep the subscript α to be consistent with the notation in the rest of the paper.

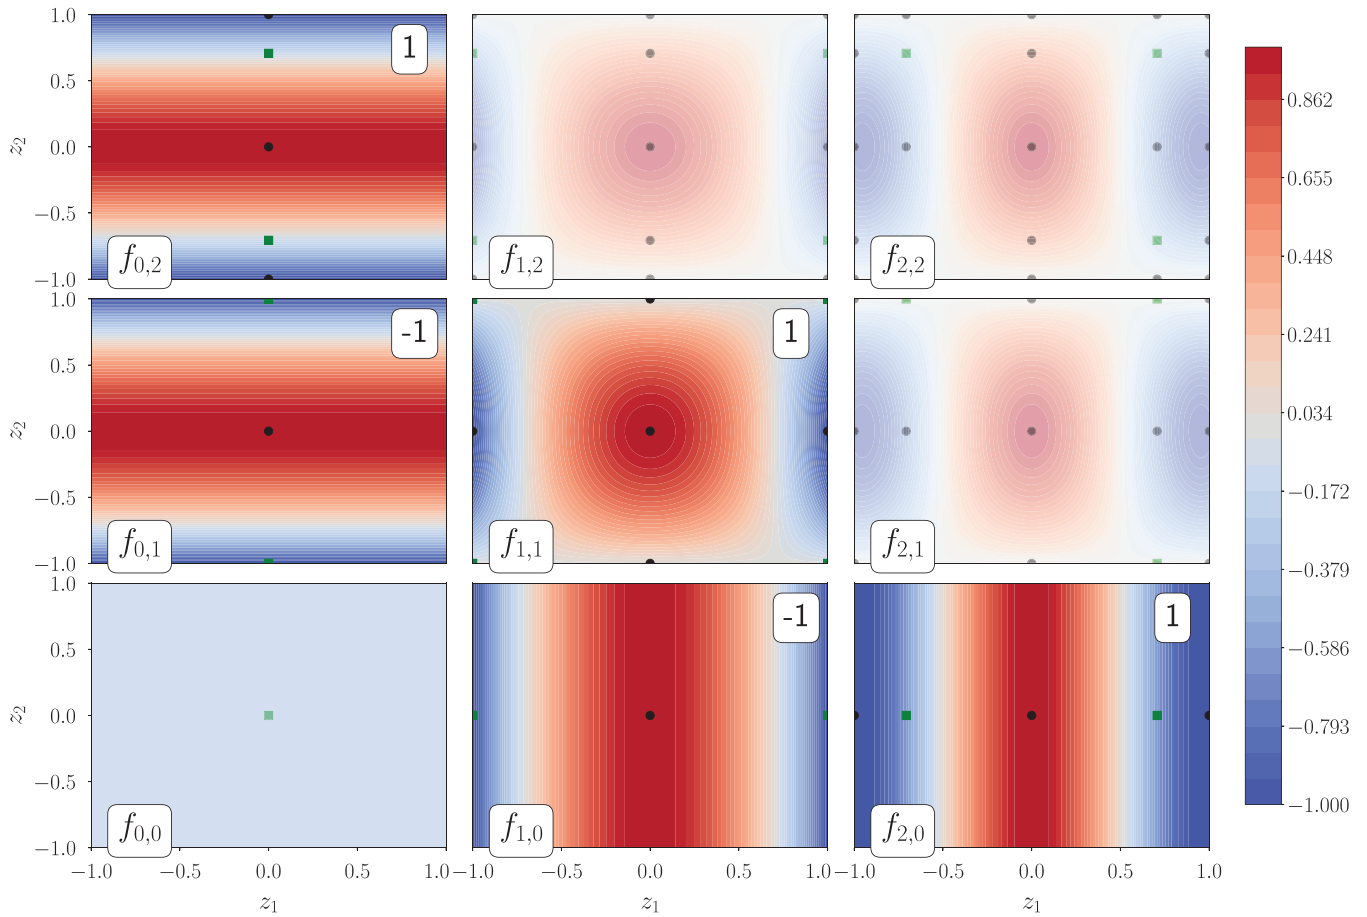


FIGURE 5 The tensor-product grids and contours of the interpolants used to construct the level $l = 2$ isotropic sparse grid $\hat{f}_{\alpha, l(2)}(\mathbf{z})$. The numbers in the top right of each plot are the coefficients c_{β} in (18). The tensor interpolants without coefficients do not contribute to the sparse interpolant. The green squares depict the unique samples $\mathbf{z}_{\beta}^{\text{diff}}$ of the tensor-product grid \mathcal{Z}_{β} that do not exist in another tensor-product grid \mathcal{Z}_{β^*} , $\beta^* \leq \beta$. Both the samples depicted by the black circles and the green squares are used to build an interpolant

interpolants, which require large numbers of points, is the primary mechanism by which sparse grids mitigate the curse of dimensionality.

Note the approximation $\hat{f}_{0,0}$ is not included in the sparse grid interpolant. As the level l is increased more lower-level tensor-product interpolants will be excluded from the sparse grid approximation.

4 | MULTI-INDEX COLLOCATION

Sparse-grids and other surrogate methods can mitigate the curse of dimensionality by exploiting structure in the function being approximated. However, many if not all surrogate methods can be infeasible in high-dimensions if the computational cost of evaluating a model is high. In this section, we present a method for building surrogates in such situations, which leverages simulations of lower-fidelity models of reduced cost to increase the tractability of building surrogates of a high-fidelity model. We focus on sparse grid function approximation; however, in Section 5, we briefly discuss how the algorithm we use for managing evaluations of models of varying fidelity can be generalized to other types of approximations, such as polynomial chaos expansion constructed on arbitrary points sets, Gaussian processes, low-rank approximations, or even neural networks.

Using the triangle inequality, it is straightforward to decompose the error in the sparse grid approximation $\hat{f}_{\alpha, l}$ of the exact model f , built using evaluations of \hat{f}_{α} , into the sum of the a deterministic and stochastic approximation errors.

$$\|f - \hat{f}_{\alpha, \mathcal{I}}\|_{L^p(\Gamma)} \leq \|f - \hat{f}_\alpha\|_{L^p(\Gamma)} + \|\hat{f}_\alpha - \hat{f}_{\alpha, \mathcal{I}}\|_{L^p(\Gamma)}. \tag{22}$$

Here, the first term on the right-hand side represents the deterministic error and the second term represents the stochastic error. These two errors should be balanced to minimize inefficiency resulting from over-resolving a nondominant source of error. Figure 2 demonstrates that mesh discretizations can effect the accuracy of statistics computed using a numerical model. For example, when the time step is to large, refining the spatial mesh has no effect on error in the mean $\mathbb{E}[\hat{f}_\alpha]$. Analogously (22) implies that refining the mesh when the stochastic error is larger will have no effect on overall error $\|f - \hat{f}_{\alpha, \mathcal{I}}\|_{L^p(\Gamma)}$.

Multi-level/Multi-index collocations methods^{22,23,26,41} combine a sequence of approximate physical models with a sequence of interpolants both to balance physical and stochastic errors and to reduce the computational cost of achieving a specified level of accuracy. In this paper, we extend the multi-index stochastic collocation method (MISC) presented in Reference 23. The other multilevel methods cited can be formulated as specialized instances of MISC.

Given an ensemble of models $\mathcal{M}_\alpha = \otimes_{i=1}^{n_\alpha} \mathcal{M}_{\alpha_i}$ of varying physical fidelity (numerical discretization), which can be used to approximate the QoI f , MISC approximates the QoI and its mean with the following generalization of Equation (14)

$$\hat{f}_{\mathcal{J}}(\mathbf{z}) = \sum_{[\alpha, \beta] \in \mathcal{J}} c_{\alpha, \beta} \hat{f}_{\alpha, \beta}(\mathbf{z}), \quad \mu_{\mathcal{J}} = \sum_{[\alpha, \beta] \in \mathcal{J}} c_{\alpha, \beta} \mu_{\alpha, \beta}, \tag{23}$$

MISC works most effectively provided the following assumption

Assumption 1. Given a sequence of tolerance parameters, the error in the successive approximations of f decreases as the fidelity increases, that is, $\|\hat{f}_{\alpha^*} - f\|_{L^p(\Gamma)} < \|\hat{f}_\alpha - f\|_{L^p(\Gamma)}$, if $\alpha^* > \alpha$.

Similarly to sparse grid function approximation (see Section 3.2.2), the structure of the index set \mathcal{J} of the MISC approximation (23) directly controls the computational cost of building, and the accuracy of, a multi-index surrogate. However, unlike the sparse grid index set \mathcal{I} , the MISC index set \mathcal{J} not only controls the allocation of samples to each random variable, but also the allocation of resources to each of the model discretizations.

Given a MISC approximation $\hat{f}_{\mathcal{J}}$, let $\hat{f}_{\mathcal{J} \cup \{[\alpha, \beta]\}}$ be the new MISC approximation obtained by including a new tensor-product interpolant $\hat{f}_{\alpha, \beta}$. Similarly to Section 3.2.2, let

$$\Delta E_{\alpha, \beta} = \|\hat{f}_{\mathcal{J} \cup \{[\alpha, \beta]\}} - \hat{f}_{\mathcal{J}}\| \quad \Delta W_{\alpha, \beta} = |\text{Work}[\hat{f}_{\mathcal{J} \cup \{[\alpha, \beta]\}}] - \text{Work}[\hat{f}_{\mathcal{J}}]|, \tag{24}$$

denote the difference between two successive MISC approximations and the work needed to update the MISC approximation, respectively. For a given α , the cost of running a simulation at a single realization of \mathbf{z} is again denoted as W_α , such that we have $\Delta W_{\alpha, \beta} = W_\alpha \text{card}(\mathcal{Z}_{\alpha, \beta}^{\text{diff}})$.

With the goal of generating a MISC approximation that balances the computational work and error, we can again formulate and solve a knapsack problem to define the index set \mathcal{J} . The knapsack problem is identical to (17), except now the optimization is not just over $\beta \in \mathbb{N}$, but also over $\alpha \in \mathbb{N}$. If the quantities $\Delta E_{\alpha, \beta}$ and $\Delta W_{\alpha, \beta}$ are available, the quasi-optimal** solution to the knapsack problem is

$$\mathcal{J} = \left\{ [\alpha, \beta] \in \mathbb{N}_0^{n_\alpha + n_\beta} \mid \frac{\Delta E_{\alpha, \beta}}{\Delta W_{\alpha, \beta}} > \epsilon \right\}, \tag{25}$$

where ϵ is chosen to reflect the desired accuracy in the MISC approximation.²³ Computing this index set is not practical as it requires computing all possible interpolants and then choosing the approximations which significantly contribute to the MISC error. The interpolants ignored would contribute little to error but may incur significant cost. In this paper, we propose using an adaptive algorithm to greedily choose the most cost effective interpolants to add to a index set \mathcal{J} . We will present our adaptive algorithm in Section 4.2; however, we first provide an example of using MISC with an isotropic index set to approximate a simple algebraic ensemble of models. This example is intended to develop intuition about the fundamental mechanisms which lead to the improved efficiency of MISC, as compared to single fidelity strategies.

**The solution is only quasi optimal because it is based upon an upper bound on the approximation error, that is $\|f - \hat{f}_{\mathcal{J}}\| \leq \sum_{[\alpha, \beta] \notin \mathcal{J}} \Delta E_{\alpha, \beta}$.

4.1 | An example of isotropic MISC approximation

The isotropic MISC approximation of a function f is

$$\hat{f}_{\mathcal{J}(l)} = \sum_{[\alpha, \beta] \in \mathcal{J}(l)} (-1)^{l - \|\alpha + \beta\|_1} \binom{n_\alpha + n_\beta - 1}{l - \|\alpha + \beta\|_1} \hat{f}_{\alpha, \beta}, \quad (26)$$

where $\mathcal{J}(l) = \{[\alpha, \beta] | (\max(0, l - 1) \leq \|\alpha + \beta\|_1 \leq l + n_\alpha + n_\beta - 2)\}$. Consider the MISC approximation of the function

$$f(z) = \cos\left(\frac{\pi}{2}\left(z_1 + \frac{4}{5}\right)\right), \quad (27)$$

using the sequence of approximations $\hat{f}_{\alpha_1} = \cos\left(\frac{\pi}{2}\left(z_1 + \frac{4}{5} + \epsilon_{\alpha_1}\right)\right)$ where $\epsilon_i > \epsilon_{i+1} \geq 0$. These approximations, shown in the left of Figure 6 for $\epsilon_{\alpha_1=0} = 1/5$, $\epsilon_{\alpha_1=1} = 1/10$ and $\epsilon_{\alpha_1=2} = 1/20$, converge to the true function as epsilon is driven to zero, that is, $\hat{f}_{\alpha_1}(z) \rightarrow f(z)$ as $\alpha_1 \rightarrow \infty$.

In Figure 7, we plot the tensor-product grids and corresponding interpolants that make up the MISC approximation. The interpolants $\hat{f}_{\alpha_1, \beta}$ converge to \hat{f}_{α_1} as β increases. However for low values of α_1 the deterministic prediction error is large. Refinement of both α_1 and β is needed for the interpolants to converge to f .

Let $\mathcal{K}_{\mathcal{J}|\alpha} = \{[\alpha^*, \beta^*] | \beta^* \in \mathcal{J}, \alpha^* = \alpha\}$ denote the set of unique stochastic multi-indices for a given deterministic refinement α . Then for each physical discretization α the MISC approximation requires evaluating the function \hat{f}_α on the points

$$\mathcal{Z}_{\mathcal{K}_{\mathcal{J}|\alpha}} = \bigcup_{\beta \in \mathcal{K}_{\mathcal{J}|\alpha}} \mathcal{Z}_\beta, \quad (28)$$

where $M_{\mathcal{K}_{\mathcal{J}|\alpha}} = \text{card}(\mathcal{Z}_{\mathcal{K}_{\mathcal{J}|\alpha}})$ is the number of samples, of the random variables, at which \hat{f}_α is evaluated. The total number of samples used by the MISC approximation, for any index set \mathcal{J} , is

$$M_{\mathcal{J}} = \sum_{\alpha \in \mathcal{L}} M_{\mathcal{K}_{\mathcal{J}|\alpha}},$$

where $\mathcal{L} = \{\alpha | \alpha \in \mathcal{J}\}$ is the set of unique physical discretizations indices.

As can be seen in Figure 7, the number of evaluations allocated to an approximate model f_{α_1} decreases as the level of fidelity α_1 increases, that is, $M_{\mathcal{K}_{\mathcal{J}|\alpha}} > M_{\mathcal{K}_{\mathcal{J}|\alpha^*}}$, if $\alpha^* > \alpha$. For example, 5 samples are allocated to evaluating \hat{f}_0 , 3 to \hat{f}_1 and 1 to \hat{f}_2 .

Despite the fact only one sample is allocated to evaluating \hat{f}_2 , the MISC approximation, shown in the right plot of Figure 6, is more accurate than any of the tensor-product interpolants that make up the MISC approximation (grids with a gray background). Understanding of this observation can be gained by viewing the discrepancies between models $\hat{f}_{i+1} - \hat{f}_i$ depicted in the middle plot of Figure 6. The magnitude of these discrepancies decreases substantially as α_1 increases.




FIGURE 6 (Left) The function f (27) and the approximations \hat{f}_{α_1} . (Middle) The discrepancies between successive approximations \hat{f}_{α_1} and \hat{f}_{α_1+1} . (Right) The most accurate (largest β_1) single fidelity interpolants f_{α_1, β_1} , used by MISC, which use only evaluations of one approximate model \hat{f}_{α_1} , $\alpha_1 = 0, 1, 2$, and the MISC approximation $\hat{f}_{\mathcal{J}}$ which combines each of these interpolants (and some less accurate interpolants) [Color figure can be viewed at wileyonlinelibrary.com]

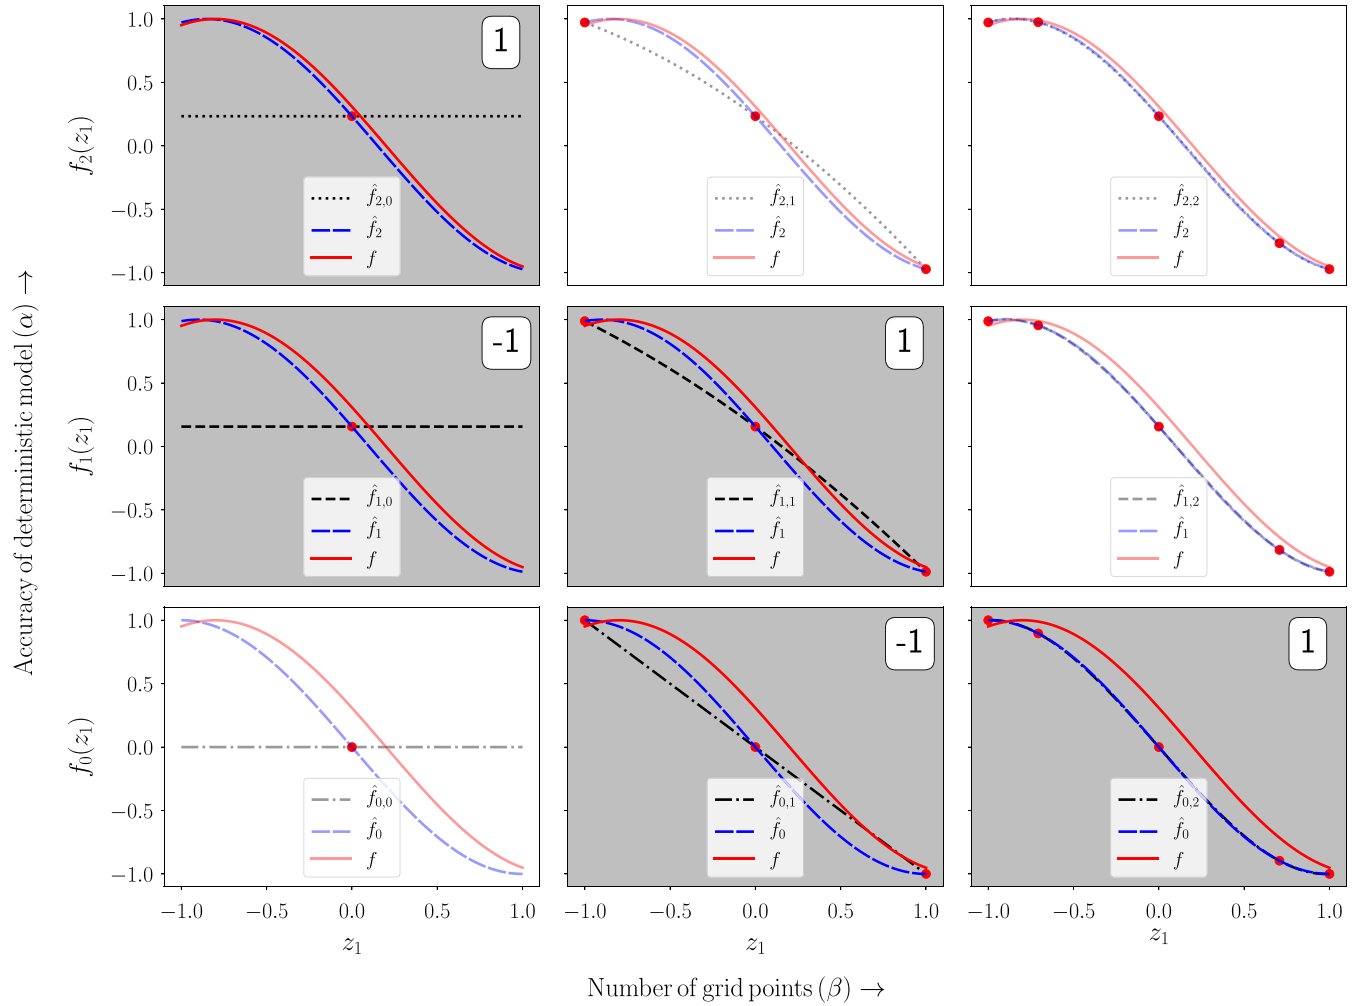


FIGURE 7 The tensor-product grids and interpolants of the level $l = 2$ isotropic MISC approximation $\hat{f}_{J(2)}(\mathbf{z})$. The grids plotted with a gray background are included in the MISC approximation, whereas the remaining grids are ignored. The numbers in the top-right of each plot are the coefficients $c_{\alpha,\beta} = (-1)^{l-\|\alpha+\beta\|_1} \binom{d-1}{l-\|\alpha+\beta\|_1}$ in (26). The dash-dotted black lines represent the tensor-product interpolants $\hat{f}_{\alpha,\beta}$. These interpolants converge to the lower fidelity models \hat{f}_α (dashed blue lines) as β increases, and to the high-fidelity model f (solid red line) as both α and β are increased [Color figure can be viewed at wileyonlinelibrary.com]

Thus a constant approximation $\hat{f}_{2,0}$ of \hat{f}_2 is sufficient to balance the deterministic and stochastic errors, that is, the error caused by nonzero values of ϵ and using a finite number of samples to build the interpolants, respectively. This constant approximation is useful because it reduces the bias in the MISC interpolant by decreasing the physical discretization error. In contrast the approximation $\hat{f}_{0,2}$ is useful as it captures the variation of the function f with respect to the variable z_1 thereby reducing the stochastic error.

We remark here that the total number of samples used to build the MISC approximation is much larger than the number of samples used to build $\hat{f}_{2,2}$. In general, the number of points in a MISC approximation is much greater than the number of points in a single fidelity sparse grid approximation of the highest fidelity model f_{α^*} , with equivalent cost, that is, $M(J) \gg M_{\alpha^*,I}$. However, for nontrivial models, the total work needed to build the MISC approximation \hat{f}_J will be less than the work needed to build the single fidelity sparse grid approximation $f_{\alpha^*,I}$ of f_{α^*} , since many of these samples occur at lower fidelity.

4.2 | Adaptivity

The isotropic MISC approximation (26) balances computational expense with accuracy. The same number of model evaluations are used for two physical discretization parameters $\alpha \neq \alpha^*$ if $\|\alpha\|_1 = \|\alpha^*\|_1$. For example the same number of

model evaluations will be used to evaluate \hat{f}_α and \hat{f}_{α^*} if $n_\alpha = 2$ and $\alpha = (2, 1)$ and $\alpha^* = (1, 2)$. The isotropic formulation does not account for differences in the cost of evaluating the two models. If \hat{f}_α is twice as expensive as evaluating \hat{f}_{α^*} , then ideally this should be reflected in the sample allocation used by the MISC approximation. The isotropic formulation also assumes each variable and \hat{f}_{α^*} contributes equally to the error in the MISC approximation. However, often only a subset of variables significantly influences the variability of a function. This parameter sensitivity should also be accounted for when building the MISC approximation.

We propose using a greedy algorithm to adaptively choose the most cost effective interpolants to add to the MISC approximation. Our adaptive algorithm is based on the algorithm proposed in Reference 42 and pseudo-code for the algorithm is shown in Algorithm 1. The algorithm is initialized with a MISC approximation comprising of a set of indices \mathcal{J} representing the current approximation, a set of active indices \mathcal{A} that indicate the physical discretization, random dimensions for potential refinement, and sparse grid coefficients $C = \{c_{\alpha,\beta}\}_{\alpha,\beta \in \mathcal{J}}$. Often $\mathcal{J} = \{\mathbf{0}\}$, $\mathcal{A} = \{\mathbf{e}_k, k \in [n_\alpha + n_\beta]\}$, and $C = \{c_{\mathbf{0},\mathbf{0}} = 1\}$. Here, we use \mathbf{e}_k to denote the unit vector with a value of 1 in the k th position.

Algorithm 1 INTERPOLATE $[\{\hat{f}_\alpha(\mathbf{z})\}, C, \mathcal{J}, \mathcal{A}, \tau, W_{\max}] \rightarrow \hat{f}_J$

```

1: while NOT TERMINATE $[\mathcal{A}, N, \tau, W_{\max}]$  do
2:    $[\alpha^*, \beta^*] := \operatorname{argmax}_{[\alpha,\beta] \in \mathcal{A}} \gamma_{\alpha,\beta}$  ▷ Determine the index with the highest priority
3:    $\mathcal{A} := \mathcal{A} \setminus [\alpha^*, \beta^*]$  ▷ Remove  $[\alpha^*, \beta^*]$  from the active set
4:    $C = \text{UPDATE}[\mathcal{J}, [\alpha^*, \beta^*], C]$  ▷ Update the sparse grid coefficients
5:    $\mathcal{J} := \mathcal{J} \cup [\alpha^*, \beta^*]$  ▷ Add  $[\alpha^*, \beta^*]$  to the MISC set
6:    $\mathcal{R} := \text{REFINE}[\alpha^*, \beta^*, \mathcal{J}]$  ▷ Find all admissible forward neighbors of  $[\alpha^*, \beta^*]$ 
7:    $\gamma_{\alpha,\beta} := \text{INDICATOR}[\alpha, \beta] \forall [\alpha, \beta] \in \mathcal{R}$  ▷ Calculate the priority of the neighbors
8:    $\mathcal{A} := \mathcal{A} \cup \mathcal{R}$  ▷ Add the forward neighbors to the active index set
9: end while

```

The algorithm begins by selecting the index $[\alpha^*, \beta^*] \in \mathcal{A}$ with the largest refinement indicator $\gamma_{\alpha,\beta}$. The interpolant $\hat{f}_{\alpha^*,\beta^*}$ is added to the sparse grid and its indices are used to identify new candidate interpolants for consideration in reducing the approximation error. This process continues until a computational budget (W_{\max}) limiting the total computational work ($\sum_{\mathcal{J}} \Delta W_{\alpha,\beta}$) is reached or until a global error indicator (eg, $\sum_{\mathcal{A}} \Delta E_{\alpha,\beta}$) drops below a predefined threshold τ . These exit criteria are checked using the TERMINATE routine in step 1 of Algorithm 1. The UPDATE function from step 4 of Algorithm 1 is summarized in Algorithm 2; it adds the interpolant $\hat{f}_{\alpha^*,\beta^*}$ to the MISC approximation. The INDICATOR routine from step 7 of Algorithm 1 controls which indices are added to the sparse grid via the use of index error and global error metrics. These indicators provide estimates of the contribution of an index to reducing the error in the interpolant, and of the error in the entire interpolant, respectively.

Algorithm 2 UPDATE $[C_{\mathcal{J}}, \mathcal{J}, [\alpha^*, \beta^*]] \rightarrow C_{\mathcal{J} \cup \{[\alpha^*, \beta^*]\}}$

```

1:  $c_{\alpha^*,\beta^*} = 1$ 
2: for  $[\alpha, \beta] \in \mathcal{J}$  do
3:    $\Delta = [\alpha, \beta] - [\alpha^*, \beta^*]$ 
4:   if  $\min_{i \in [n_\alpha + n_\beta]} \Delta \geq 0$  and  $\max_{i \in [n_\alpha + n_\beta]} \Delta \leq 1$  then
5:      $c_{\alpha,\beta} = c_{\alpha,\beta} + (-1)^{\|\Delta\|_1}$ 
6:   end if
7: end for

```

In this paper we use the following indicators to guide refinement

$$\begin{aligned}
\Delta E_{\alpha,\beta}^\mu &= \frac{1}{|\hat{f}_{\mathbf{0},\mathbf{0}}|} |\mathbb{E}[\hat{f}_{\mathcal{J} \cup \{[\alpha,\beta]\}}(\mathbf{z})] - \mathbb{E}[\hat{f}_{\mathcal{J}}(\mathbf{z})]| & \Delta E_{\alpha,\beta}^{\sigma^2} &= \frac{1}{|\hat{f}_{\mathbf{0},\mathbf{0}}|^2} |\mathbb{V}[\hat{f}_{\mathcal{J} \cup \{[\alpha,\beta]\}}(\mathbf{z})] - \mathbb{V}[\hat{f}_{\mathcal{J}}(\mathbf{z})]| \\
\gamma_{\alpha,\beta} &= \frac{1}{\Delta W_{\alpha,\beta}} (\kappa \Delta E_{\alpha,\beta}^\mu + (1 - \kappa) \Delta E_{\alpha,\beta}^{\sigma^2}), & \kappa &\in [0, 1].
\end{aligned} \tag{29}$$




FIGURE 8 Two steps of the adaptive sparse grid algorithm for $n_\alpha = n_\beta = 1$. The gray and red boxes represent the sets \mathcal{J} and \mathcal{A} , respectively. The numbers within the boxes are the sparse grid coefficients. The second and fourth plots represent the samples assigned to each level of physical discretization level α_1 . The points corresponding to the tensor-product interpolations indexed by the set \mathcal{J} and the active set \mathcal{A} are given in black and red, respectively [Color figure can be viewed at wileyonlinelibrary.com]

The refinement indicator $\gamma_{\alpha,\beta}$ attempts to maximize the reduction of the error in the mean and variance of the approximation whilst minimizing cost. Specifically the error indicator measures (via a convex combination) the contribution of the interpolant $\hat{f}_{\alpha,\beta}$ to the mean and variance of the MISC approximation $\hat{f}_{\mathcal{J}}$ which can be computed using (23).

We remark here that initialization challenges exist when using $\kappa = 0$ to construct an adaptive MISC approximation. For any index $[\alpha, \beta]$ with $\beta = \mathbf{0}$, the interpolants $\hat{f}_{\alpha,\beta}$ are built using only a single point and so are just constant functions. For these interpolants, the variance is zero and so not a useful indicator of the impact of refinement for initial α candidates. Also note that the error indicators, $\Delta E_{\alpha,\beta}^\mu$ and $\Delta E_{\alpha,\beta}^{\sigma^2}$, are relative to the magnitude of the function and its square at the center of the sparse grid. This is important to ensure that mean or the variance components of the indicator dominate simply because they are on a different scales.

The index $[\alpha, \beta]$ with the largest refinement indicator is refined using the function REFINE (step 6 of Algorithm 1) by adding all indices $[\alpha, \beta] + \mathbf{e}_k, k \in [n_\alpha + n_\beta]$ that satisfy the following downward-closed admissibility criterion

$$[\alpha, \beta] + \mathbf{e}_k - \mathbf{e}_j \in \mathcal{J} \text{ for } j \in [n_\alpha + n_\beta], l_k > 1. \tag{30}$$

The active set \mathcal{A} is then rebuilt by adding each index corresponding to the indices from (30).

Figure 8 depicts two steps of the adaptive sparse grid algorithm for $n_\alpha = n_\beta = 1$. The first and third plots depict the index sets \mathcal{J} and \mathcal{A} represented by the gray and red, respectively. The numbers within the boxes are the sparse grid coefficients. The second and fourth plots represent the samples assigned to each level of physical discretization level $\alpha_1, l_{\alpha_1} = 0, \dots, 3$. The points corresponding to the tensor-product interpolants indexed by the set \mathcal{J} and the active set \mathcal{A} are given in black and red, respectively. The new work required to refine an index is given by the cost of running the simulation at all the red points $\mathcal{Z}_{\alpha,\beta}^{\text{diff}}$. In a previous step, the index $(\alpha_1, \beta_1) = (0, 1)$ was chosen for refinement. Both forward neighbors $\mathbf{k} = [\alpha, \beta] + \mathbf{e}_n, n \in [n_\alpha + n_\beta]$ are added to the active set \mathcal{A} (red boxes in first figure) because these indices satisfy the admissibility criteria. This requires evaluating both the functions \hat{f}_0 and \hat{f}_1 at 2 new points. In the next step, the index $(\alpha_1, \beta_1) = (0, 2)$ is selected for refinement. In this step, only one new index $(\alpha_1, \beta_1) = (0, 2) + (0, 1) = (0, 3)$ is added to the active set (third figure) because the other forward neighbor $(\alpha_1, \beta_1) = (0, 2) + (1, 0) = (1, 2)$ does not satisfy the admissibility criteria. Adding the index $(2, 0)$ to \mathcal{A} requires evaluating \hat{f}_0 at 4 new points.

4.3 | Multilevel collocation

Multilevel collocation (MLC)²² is a specialized instance of multi-index collocation. MLC assumes that there is only one discretization index which controls the model accuracy and cost, that is, $\alpha = (\alpha_1)$. We can easily apply the AMISC algorithm to the multilevel setting by simply defining the index set \mathcal{M}_α correctly. For the advection-diffusion considered in this paper we make the reasonable choice that $\mathcal{M}_\alpha = \{\otimes_{i=1}^3 \{1\}, \dots, \otimes_{i=1}^3 \{k\}\}, k \in \{1, \dots, 5\}$. This multi-index model set has a one-to-one mapping with the 1D multilevel hierarchy $\hat{\mathcal{M}}_\alpha = \{1, 2, 3, 4, 5\}$. We do not consider multi-level methods for the nozzle problem as it is unclear how to define a reasonable 1D hierarchy for that model.

The original multi-level algorithm²² allocate samples to models of different model discretizations using a priori theoretical error estimates. Consequently it is challenging to apply this approach to practical problems where such theoretical estimates are unavailable. A partially adaptive a posteriori MLC strategy was proposed in Reference 43. However, when

using this approach, the resolution of the approximation built over the stochastic space is a priori tied to the resolution of the numerical solver used to solve the PDE. The adaptive multi-index collocation presented here can easily be applied in the multi-level setting without the restrictions of the aforementioned approaches. The computational cost and accuracy of multi-level and multi-index sparse grid collocation are compared in Section 6.1.

Finally it is important to note that multi-level surrogate methods have been developed which do not use sparse grids to build multivariate approximations over the stochastic space. Gaussian processes were used in References 18-20 and low-rank reduced order models were used in Reference 21. These methods are not considered further here since, similar to multi-level sparse grid collocation, if a model does indeed have multiple discretization hyper-parameters it can be unclear how to define a suitable 1D model hierarchy.

5 | MULTI-INDEX SENSITIVITY ANALYSIS

Quantifying the sensitivity of a model output f to the model parameters \mathbf{Z} can be an important component of assessing uncertainty in model predictions. For example, parameter sensitivities can be used to prioritize data collection. Using data or additional knowledge to reduce uncertainty in sensitive model inputs can significantly decrease prediction uncertainty. This section demonstrates how to use MISC for global sensitivity analysis, specifically for computing Sobol indices.⁴⁴

Sobol indices are frequently used to estimate the sensitivity of a function to single, or combinations of input parameters. Sobol indices can be estimated numerically using Monte Carlo methods;⁴⁵ however, they can be computed analytically from a sparse grid.^{46,47} For any tensor-product probability density function ω , there exists a transformation of the Lagrange polynomial basis into a orthonormal polynomial basis,⁴⁸ also known as a polynomial chaos expansion (PCE). Applying such a transformation facilitates simple and analytical computation of Sobol indices from the coefficients of the PCE.² The following sections describe the basic properties of a PCE and how to transform an MISC approximation into a PCE for the purpose of computing Sobol indices.

5.1 | Polynomial chaos expansions

Polynomial chaos expansions (PCE)^{1,3,49} represent the model output $\hat{f}_\alpha(\mathbf{z})$ as an expansion of orthonormal polynomials

$$\hat{f}_\alpha(\mathbf{z}) \approx \hat{f}_{\alpha,\Lambda}(\mathbf{z}) = \sum_{\lambda \in \Lambda} \eta_\lambda \phi_\lambda(\mathbf{z}), \quad |\Lambda| = N. \quad (31)$$

The basis functions ϕ_λ are typically constructed to be orthonormal with respect to the density ω , that is

$$(\phi_{\lambda^{(i)}}(\mathbf{z}), \phi_{\lambda^{(j)}}(\mathbf{z}))_{L^2_\omega(\Omega)} := \int_\Gamma \phi_{\lambda^{(i)}}(\mathbf{z}) \phi_{\lambda^{(j)}}(\mathbf{z}) d\omega(\mathbf{z}) = \delta_{i,j},$$

where $\delta_{i,j}$ is the Kronecker delta function and $\lambda^{(i)}, \lambda^{(j)}$ are two different multivariate indices where $\lambda = (\lambda_1 \dots, \lambda_d) \in \mathbb{N}_0^d$. When the components of \mathbf{z} are independent we can express the basis functions ϕ as tensor products of univariate orthonormal polynomials. That is

$$\phi_\lambda(\mathbf{z}) = \prod_{i=1}^d \phi_{\lambda_i}^i(z_i), \quad \int_{\Omega_i} \phi_j^i(z_i) \phi_k^i(z_i) \omega_i(z_i) dz_i = \delta_{j,k}, \quad j, k \geq 0, \quad \deg \phi_j^i = j. \quad (32)$$

The mean and variance of the PCE can be computed directly from the expansion coefficients.

$$\mathbb{E}[\hat{f}_{\alpha,\Lambda}] = \eta_0, \quad \mathbb{V}[\hat{f}_{\alpha,\Lambda}] = \sum_{\substack{\lambda \in \Lambda \\ \lambda \neq 0}} \eta_\lambda^2. \quad (33)$$

The Sobol indices can also be computed analytically from the expansion using the identity

$$\mathbb{V}[\hat{f}_\alpha] = \sum_{\lambda \in \Lambda_\alpha} \eta_\lambda^2 \quad \Lambda_\alpha = \{\lambda | \lambda_i > 0, i \in \alpha, \lambda_j = 0, j \notin \alpha\}. \quad (34)$$

5.2 | Transforming a sparse grid into a polynomial chaos expansion

Let \mathcal{T} denote the transformation of a Lagrange polynomial interpolant $\hat{f}_{\alpha,\beta}$ to an orthonormal polynomial interpolant $\hat{f}_{\alpha,\Lambda}$. From (23), we have $\mathcal{T}[\hat{f}_J(\mathbf{z})] = \sum_{[\alpha,\beta] \in \mathcal{J}} c_{\alpha,\beta} \mathcal{T}[\hat{f}_{\alpha,\beta}(\mathbf{z})]$, where c_β are the Smolyak coefficients. Thus, to transform the sparse grid, we only need the ability to transform each tensor-product interpolation individually. To apply the transformation of the tensor-product interpolants, we must only compute transformations of 1D basis functions. In one dimension, the polynomial interpolant of a set of points is unique. Consequently, we can express the j -th univariate Lagrange basis function in the i -th variable dimension as

$$l_{i,j}(z_i) = \sum_{k=1}^{m_{\beta_i}} v_k \phi_k^i(z_i) \quad v_k = \int_{\Gamma_k} l_{i,j}(z_i) \phi_k^i(z_i) d\omega(z_i) \quad j = 1, \dots, m_{\beta_i}, i = 1, \dots, d. \tag{35}$$

We can compute the coefficients of the orthonormal basis using Gaussian quadrature. Note this step does not require any evaluations of the model \hat{f}_α and can be precomputed offline and stored for future use. Given the orthonormal representation of the univariate Lagrange polynomials, the tensor-product interpolant can be expressed as

$$\hat{f}_{\alpha,\beta}(\mathbf{z}) = \sum_{j \leq \beta} \hat{f}_\alpha(\mathbf{z}^{(j)}) \prod_{i=1}^d l_{i,j_i}(z_i) = \sum_{j \leq \beta} \hat{f}_\alpha(\mathbf{z}^{(j)}) \prod_{i=1}^d \left(\sum_{k=1}^{m_{\beta_i}} v_k \phi_k^i(z_i) \right) = \sum_{\lambda \in \Lambda_\beta} \eta_\lambda \phi_\lambda(\mathbf{z}), \tag{36}$$

where $\Lambda_\beta = \{\lambda | \lambda \leq \beta\}$. This expression can be now used with (34) to compute Sobol indices analytically and efficiently from the sparse grid.

Note that the algorithm described here is different from the algorithms proposed in References 46 and 47. These approaches solve a multivariate interpolation problem to compute the coefficients γ of the tensor-product PCE basis in (36). The dominant cost of that approach is the inversion of the multivariate Vandermonde-like matrix evaluated at each of the tensor product points, which requires $O(M^3(\beta))$ operations. By focusing on univariate transformations of the Lagrange basis, we can compute the coefficients of the multivariate PCE via an outer product of the univariate orthonormal basis coefficients at a cost of $O(M^2(\beta))$. This reduction in the cost of computing the PCE transformation is practically significant, but the complexity of the sparse grid transformation still grows linearly with the number of points in the sparse grid. Note that the cost of either transformation algorithm is typically negligible to the cost of evaluating the models \hat{f}_α .

Remark 1. In this paper, we use sparse grids to form a MISC approximation. Sparse grids naturally combine the refinement of the discretization parameters and the refinement of the interpolants. Noting that the MISC approximation can be expressed in the general form $\hat{f}_J(\mathbf{z}) = \sum_{[\alpha,\beta] \in \mathcal{J}} c_{\alpha,\beta} \hat{f}_{\alpha,\beta}(\mathbf{z})$, Equation (36) suggest that sparse grids are not the only means to build a MISC approximation. Specifically, one could draw from a broad class of surrogate models to construct the approximations $\hat{f}_{\alpha,\beta}$. One could use PCE constructed using interpolation, least-squares or ℓ_1 minimization, but one could also use Gaussian processes, low-rank approximations, or even neural networks. Given a method to build $\hat{f}_{\alpha,\beta}$, the adaptive algorithm described in Section 4.2 can be used to guide the selection of the indices $[\alpha, \beta]$ so long as a sampling method, which can be controlled by the stochastic discretization parameter β , is available. It is straightforward to construct such sampling methods; for example, one approach is to let β be a scalar controlling the number of Monte-Carlo or Latin-hypercube samples used to build a regression-based PCE or Gaussian process. Alternatively, one can use sampling schemes such as multivariate Leja sequences for polynomial interpolation, which are more amenable to dimension-adaptivity.⁴⁹

6 | RESULTS

In this section, we explore the efficacy and properties of our proposed approach, first using the advection-diffusion model problem and then using the engineering model of a jet engine nozzle.

To measure the performance of an approximation, we will use the ℓ^∞ error on a set of test nodes. We generate a set of Q random samples $\{\mathbf{z}^{(j)}\}_{j=1}^Q \subset \Gamma$ drawn from the density ω of the uncertain variables \mathbf{Z} . The relative error is computed as

$$\|f - \hat{f}_J\|_{L^\infty(\Gamma)} = \frac{1}{\max_{j \in [Q]} |f(\mathbf{z}^{(j)})| - \min_{j \in [Q]} |f(\mathbf{z}^{(j)})|} \operatorname{argmax}_{j \in [Q]} |f(\mathbf{z}^{(j)}) - \hat{f}_J(\mathbf{z}^{(j)})|,$$

where f is the exact function and \hat{f}_J is the interpolative approximation. We set $Q = 1000$ and $Q = 500$ for the advection diffusion and nozzle models, respectively.

We use the ℓ^∞ norm to quantify error as we are interested in probabilities of failure, and convergence almost surely is required to guarantee convergence of probability density functions (PDFs) of the QoI.⁵⁰ We use the range of the values of the reference function f over the validation samples to normalize the error, enabling us to highlight how our adaptive approximations balance error among multiple QoI.

6.1 | Advection diffusion model

In this section, we will use the adaptive MISC algorithm to approximate the quantity of interest (8) obtained from the advection diffusion model (5). For this model, an efficient and accurate approximation must balance spatial refinement, temporal refinement, and the stochastic interpolation error. The details of the physical discretization are provided in Section 3.1.1.

In Figure 9, we compare the adaptive MISC algorithm (AMISC) with a single level adaptive sparse grid (see Section 3.2.2). Specifically, we plot the error in each approximation as the total work is increased. Work is measured relative to the cost of performing a single simulation using the approximate model $\hat{f}_{5,5,5}$ with the highest fidelity discretization parameters. For example, if τ_α is the wall time needed to simulate the model f_α at one realization of the parameters \mathbf{z} , then the work needed to run this model is $W_\alpha = \tau_\alpha / \tau_{5,5,5} \leq 1$. For this particular advection-diffusion problem, we do not know the true QoI exactly, so we use the model with $\alpha = (6, 6, 6)$ as the truth when computing errors. Each time the error is measured, all points associated with indices in both the active index set \mathcal{A} and the current index set \mathcal{I} are used and thus contribute to the value of work on the horizontal axis. The notation in the legend is used to denote the index sets \mathcal{M}_α defining the ensemble of models considered by the different algorithms. $\mathcal{M}_\alpha = \otimes_{i=1}^3 \{5\}$ denotes the single fidelity approximation using evaluations of only the highest fidelity model $\hat{f}_{(5,5,5)}$, $\mathcal{M}_\alpha = \otimes_{i=1}^3 \{1, \dots, k\}$ and $\mathcal{M}_\alpha = \{\otimes_{i=1}^3 \{1\}, \dots, \otimes_{i=1}^3 \{k\}\}$, respectively denote multi-index and multi-level approximation where the highest fidelity model evaluated is $\hat{f}_{(k,k,k)}$, $k \in \{3, 4, 5\}$. The gray boxes refer to two iterations during the evolution of the MISC algorithm. The point labeled 1 is taken when 500 model evaluations (of any fidelity) have been run and label 2 after 2000 model evaluations. The cost profiles of the MISC algorithm at these two iterations are depicted in Figures 10 and 11.

The error in both single-level sparse grid and AMISC approximations eventually saturate when the physical discretization error of the highest fidelity approximate model is reached. As the maximum value of α_1 is increased, the error in the approximations saturates at a smaller error. The AMISC approximations, however, achieve the same level of error as the single fidelity approach at only a fraction of the work. The results in Figure 9 are generated by setting $\kappa = 1$ in (29). Despite the fact we only adapt to reduce the error in the mean, we still can reduce the ℓ^∞ error in the AMISC and single-fidelity approximations.




FIGURE 9 (Left) Comparison of the convergence of single-fidelity and multi-fidelity approximations of the advection-diffusion model. (Right) Comparison of AMISC with adaptive multilevel approximation. The legend labels denote the index sets \mathcal{M}_α defining the ensemble of models considered by the different algorithms. $\mathcal{M}_\alpha = \otimes_{i=1}^3 \{5\}$ denotes the single fidelity approximation using evaluations of only the highest fidelity model $\hat{f}_{(5,5,5)}$, $\mathcal{M}_\alpha = \otimes_{i=1}^3 \{1, \dots, k\}$ and $\mathcal{M}_\alpha = \{\otimes_{i=1}^3 \{1\}, \dots, \otimes_{i=1}^3 \{k\}\}$ respectively denote multi-index and multi-level approximation where the highest fidelity model evaluated is $\hat{f}_{(k,k,k)}$, $k \in \{3, 4, 5\}$




FIGURE 10 The number of simulations (left) and fraction of total work (right) assigned to each approximate advection diffusion model at the point labeled 1 in Figure 9. Here $M_{\mathcal{K}_{J|\alpha}} = \text{card}(\mathcal{Z}_{\mathcal{K}_{J|\alpha}})$ is the number of samples, of the random variables, at which \hat{J}_α is evaluated. The number on top of the bars in the left plot are the exact values of $M_{\mathcal{K}_{J|\alpha}}$ and the numbers on the bars on the right plot are the fraction of total work assigned to each model indexed by α . Finally the labels of the horizontal axis represent the indices α to which AMISC has assigned at least one evaluation [Color figure can be viewed at wileyonlinelibrary.com]




FIGURE 11 The number of simulations (left) and fraction of total work (right) assigned to each approximate advection diffusion model at the point labeled 2 in Figure 9. Here $M_{\mathcal{K}_{J|\alpha}} = \text{card}(\mathcal{Z}_{\mathcal{K}_{J|\alpha}})$ is the number of samples, of the random variables, at which \hat{J}_α is evaluated. The number on top of the bars in the left plot are the exact values of $M_{\mathcal{K}_{J|\alpha}}$ and the numbers on the bars on the right plot are the fraction of total work assigned to each model indexed by α . Finally the labels of the horizontal axis represent the indices α to which AMISC has assigned at least one evaluation [Color figure can be viewed at wileyonlinelibrary.com]

The AMISC algorithm reduces the total computational cost by over two orders of magnitude, except in regions where the physical discretization error of the highest fidelity model begins to dominate.^{††} This reduction is achieved by sampling the lower-level model approximations more than the higher-level models, and only sampling the high-fidelity models when the physical discretization error starts to dominate the stochastic error. Figures 10 and 11 plot the total number of

^{††}If the cost profile of a set of model resolutions increases more rapidly, then the computational benefit of AMISC also increases. For example, if the conservative cost model for solving the advection-diffusion problem presented in Section 3.1.1 is replaced with a cost model that grows more quickly with the mesh and temporal resolution, then the benefits of the multi-fidelity methods will increase relative to their single fidelity counterparts.

simulations the MISC algorithm assigns to each approximate model for two points in the algorithm's evolution (labeled 1 and 2 in Figure 9). When the stochastic error is approximately equal to the saturation point of the blue curve, most simulations are allocated to the low-fidelity models which only involve models with indices $\alpha_i \leq 3$, $i = 1, \dots, 3$ (see Figure 10). However when the stochastic error is below the black curve, numerous samples are allocated to the higher-fidelity models with indices $\alpha_i > 3$ (see Figure 11). Although most evaluations are assigned to the lowest-level models, the highest fidelity models contribute significantly to the total work (Figure 11 right) despite being evaluated much less frequently. This is the behavior we desire. Specifically, the high-fidelity models are only evaluated when needed to maintain deterministic prediction accuracy.

The right plot of Figure 9 compares multi-index collocation with multi-level collocation using the multi-level hierarchy defined in Section 4.3. Multi-index collocation obtains significant computational savings when high-levels of accuracy are required. The computational savings increase as the number of levels utilized by the multi-level method increase. Furthermore, if the cost differential across a set of model resolutions becomes more significant, the computational benefit of MISC compared to MLC will increase.

The results generated using MLC were found using, what turned out to be, a reasonable 1D model hierarchy. It is easy to construct other hierarchies which significantly degrade the performance of MLC relative to AMISC. One major benefit of AMISC is that when a model does have multiple discretization hyper-parameters, one does not need to decide how to form a 1D model hierarchy.

As expected, the multi-level and multi-index approximation saturate at the same points when fixing the highest fidelity model used in the approximations consistently. The multi-index method does take longer to reach the exact saturation point, however, the error gets close to the saturation point much more quickly than the multi-level method. The small inefficiencies of the multi-index near the saturation points are caused by the downward-closed index set constraint for cases where final reductions in deterministic discretization error could only be addressed using the most resolved model, resulting in some loss in relative benefits for the tails of the convergence histories. This effect can be avoided by increasing the resolution of the highest fidelity model. Doing so will not effect the computational cost of the AMISC algorithm unless the evaluations of the higher fidelity model is needed to reduce the deterministic prediction error. Future work will leverage this fact to explore extensions to AMISC approaches that do not require the highest fidelity model to be specified a priori.

6.2 | Aero-thermal-structural analysis of a jet engine nozzle

In this section, we will use the adaptive MISC algorithm to approximate the four quantities of interest obtained from the nozzle model, and then use this approximation to quantify sensitivities and uncertainty in these QoI. The four model QoI we consider are mass, thrust, load-layer temperature failure ratio, and thermal-layer failure criteria. For both failure-based QoI, a value greater than one indicates failure, and both are computed using the modified P-norm function (PN-function), which aggregates pointwise values into a global quantity:

$$g(c_i) = \left(\frac{1}{N} \sum_i^N c_i^p \right)^{\frac{1}{p}}. \quad (37)$$

As p increases, the PN-function approaches the value $\max_i(c_i)$, where c_i denote local stress or temperature values. Here, we set $p = 10$.

6.2.1 | AMISC approximation

In this section, we detail the performance of the AMISC method for approximating the aforementioned QoI of the nozzle model. When considering our nozzle model, an efficient and accurate approximation must balance CFD and structural model errors. The relative computational cost of each nozzle model approximation is given in Table 1, where cost is reported in wall time using 4 CPUs for the CFD model and 1 CPU for the structural model.

For this problem, we are interested in approximating four QoI, so we must adjust slightly the refinement indicator introduced in Section 4.2. Specifically, we use a worst case error indicator that refines the multi-index set in the AMISC approximation according to the largest error across the set of QoI. Letting $\gamma_{\alpha,\beta}^i$ denote the error indicator (29) for the i th

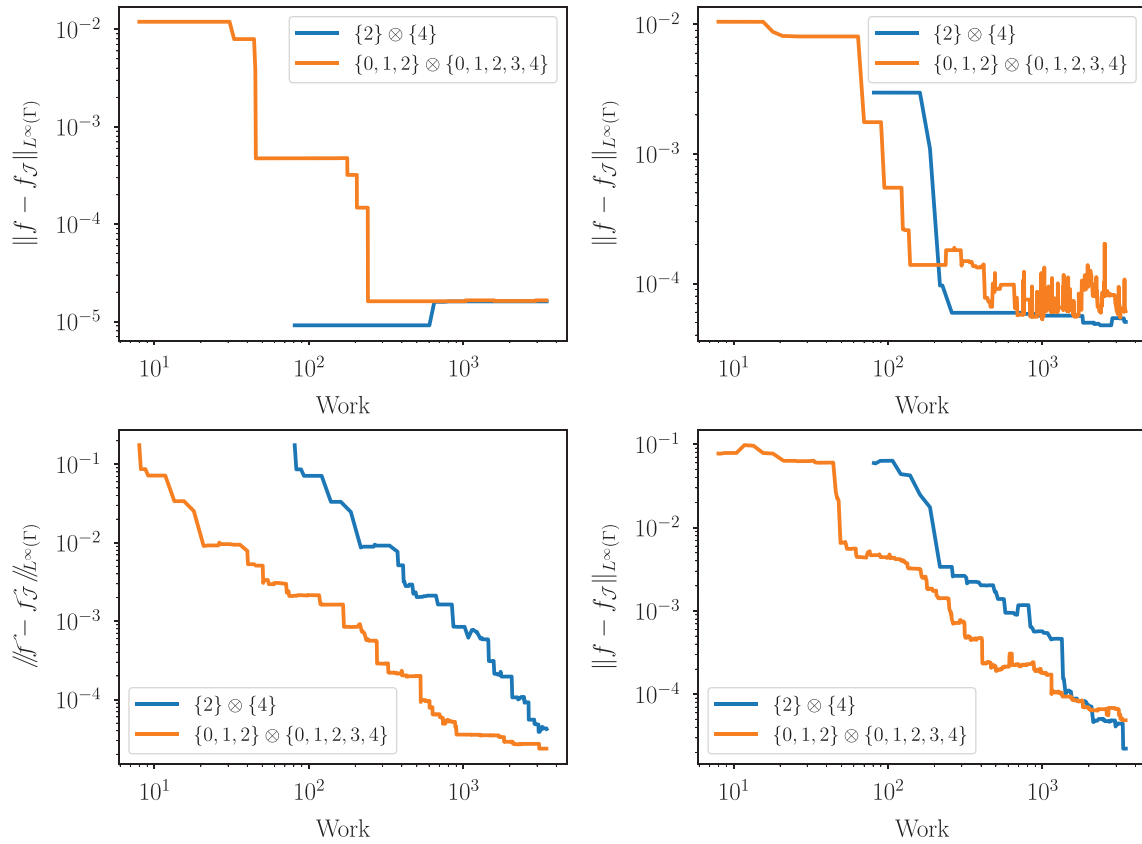


FIGURE 12 Comparison of the convergence of single-fidelity and multi-fidelity approximations of (clockwise from top left) mass, thrust, temperature layer failure and load layer failure. The legend labels denote the index sets \mathcal{M}_α defining the ensemble of models considered by the different algorithms. $\mathcal{M}_\alpha = \{2\} \otimes \{4\}$ denotes the single fidelity approximation using evaluations of only the highest fidelity model $\hat{f}_{(2,4)}$, and $\mathcal{M}_\alpha = \{0, 1, 2\} \otimes \{0, 1, 2, 3, 4\}$ denotes the multi-index approximation where the highest fidelity model evaluated is $\hat{f}_{(2,4)}$

QoI, we use

$$\gamma_{\alpha,\beta} = \max_{i \in [4]} \gamma_{\alpha,\beta}^i, \tag{38}$$

across the set of 4 QoI. We select $\kappa = 1/2$ in order to balance error reduction for the mean and variance of the approximation. The latter is consistent with our desire to use variance-based decomposition for global sensitivity analysis. The former ensures that the refinement indicator is not zero for constant functions, including approximations for $\beta = \mathbf{0}$ (see the comment in Section 4.2). Also note that the use of a relative error indicator is important to ensure that the thrust does not dominate refinement since its scale is orders of magnitude larger than the other QoI (Figure 14 shows the differing scales for the four QoI).

In Figure 12, we compare the adaptive MISC algorithm (AMISC) with a single-level adaptive sparse grid (see Section 3.2.2). Work is measured relative to the cost of simulation using the highest fidelity discretization parameters (model $\hat{f}_{3,5}$). We terminate refinement of each approximation after the cost of 3400 high-fidelity samples is reached. We again observe that, with the exception of the mass QoI, AMISC can approximate each QoI with the same accuracy as a single fidelity approximation at only a fraction of the cost. The AMISC approximation at some points on the convergence curves is an order of magnitude less expensive. This approaches the limit of potential acceleration, as there is only a factor of 20 difference in cost between the cheapest and most expensive models.

From performing parameter sweeps of each QoI using our MISC approximation we observed that mass is very close to a linear function of the random variables. Consequently, a level 1 single-fidelity sparse grid with 81 points can accurately represent mass (top left). 81 is the minimum number of points required to start the adaptive algorithm (Algorithm 1). The 81 points are used to compute the error indicators $\gamma_{\alpha,\beta}^i$ for the lowest fidelity model, indexed by $\alpha = (0, 0)$. These initial

points are comprised of one point at the center of the domain and 2 points at the upper and lower bounds of each of the 40 random parameters.

The AMISC approximation of mass achieves the same final error as the single fidelity approximation, but it is slightly less accurate for moderate cost (up to 600 equivalent high-fidelity samples). However, it is more flexible than the single-fidelity approach in that AMISC can produce reasonably accurate approximations using less than 10 equivalent high-fidelity samples. This would be important if our computational budget was more restricted than the 3400 high-fidelity samples used here.

Note that the convergence curves of the approximation of thrust (top right) is noisy in the right tail. This is related to the tolerance we use for the nonlinear CFD solver; in particular, the error is saturating at the error dictated by this tolerance. We could also include the solver tolerance as a physical discretization parameter, but this is left for future work.

The AMISC approximation of the thermal-layer failure criteria (lower right) is cheaper to construct for accuracies up to 10^{-4} . For higher accuracy, the single-fidelity approximation is cheaper to construct. In the latter situation, the two coarsest structural meshes cannot be used to reduce the total error, which is smaller than the deterministic prediction error of those models. Consequently, only evaluating the highest fidelity structural model can reduce the error. Due to the assumption that the sparse grid index set \mathcal{J} is downward closed, the AMISC procedure requires that, for a given value of the discretization parameter $\alpha_2 = k$, the models parameterized by all $\alpha_2 = j < k$ must be evaluated before the model indexed by k is evaluated. This results in redundant computation, in the form of low-fidelity evaluations when only the model $\alpha_2 = k$ can be used to reduce the error of the approximation. For this same reason, the single-level sparse grid approximation of the load-layer failure criteria (bottom left) will eventually become more efficient than the AMISC approximation if we were to increase the number of samples used to generate the convergence curves.

Note that for any amount of work, the adaptive strategy we use to build both the AMISC and single-fidelity approximations attempts to balance the error in each QoI. This is because we use a worst-case error indicator which refines the approximation so that the largest error among QoI is reduced. Over the entire evolution of the adaptive algorithm, the refinement of both the single-fidelity and AMISC approximations is primarily driven by the need to reduce error in the temperature layer criteria, since the error in the approximation of this QoI is much larger than the error in the other QoI for a given cost. For example, the error of the AMISC approximation of the temperature-layer failure criteria is approximately 4×10^{-2} when work is 100, whereas the error in the load-layer QoI is 2×10^{-3} and the errors in the mass and thrust are even smaller. Consequently, at this point in the sparse grid evolution, work will be assigned to reduce the error in the load layer criteria until its error becomes comparable to the errors in the other QoI.

In Figure 13, we plot the number of model evaluations allocated to each model discretization. The AMISC algorithm determined that the CFD and structural model are both large sources of deterministic prediction error. This is reflected by




FIGURE 13 The number of simulations (left) and fraction of total work (right) assigned to each nozzle model. Here $M_{\mathcal{K}_{\mathcal{J}|\alpha}} = \text{card}(\mathcal{Z}_{\mathcal{K}_{\mathcal{J}|\alpha}})$ is the number of samples, of the random variables, at which \hat{f}_{α} is evaluated. The number on top of the bars in the left plot are the exact values of $M_{\mathcal{K}_{\mathcal{J}|\alpha}}$ and the numbers on the bars on the right plot are the fraction of total work assigned to each model indexed by α . Finally the labels of the horizontal axis represent the indices α to which AMISC has assigned at least one evaluation [Color figure can be viewed at wileyonlinelibrary.com]




FIGURE 14 PDFs for (left to right) mass, thrust, load layer failure and temperature layer failure [Color figure can be viewed at wileyonlinelibrary.com]

the large number of samples allocated to the most refined CFD and structural models. It is interesting to note, however, that it is not necessary to refine both the CFD and structural meshes simultaneously. There is only one sample allocated to a model, $\alpha = (1, 1)$, for which both the CFD and the structural mesh have been refined. The error in the CFD model output, which is used as input to the structural model, does not significantly affect the prediction of the structural quantities of interest—the load layer and temperature layer failure criteria. From assessing the values of $\Delta E_{\alpha,\beta}^i / \Delta W_{\alpha,\beta}$ in the final grid, we determined that refinement of the CFD discretization parameter was driven mainly by the error in the mass and thrust. And not surprisingly, refinement of the structural discretization parameter was primarily driven by errors in the two structural QoI.

6.2.2 | Uncertainty quantification

In this section, we use the AMISC approximation to quantify uncertainty in the nozzle model. Specifically we report sensitivity metrics and construct PDFs for each QoI.

In Figure 14, we plot the marginal PDFs of each quantity of interest. The PDFs are approximated using Gaussian kernel density estimation using 10^6 approximate values of the QoI obtained by randomly sampling the AMISC approximation according to the distribution of the random parameters. The PDFs of both mass and thrust appear uniform. This suggests that, because the inputs are uniform, the input-output map for these QoI are primarily linear. This conclusion is also supported by plots of 1D cross sections of the map, which appear linear visually (plots omitted for brevity). The PDFs of both failure-based QoI concentrate mass on low values for each response. Since we deem failure to occur when the failure criteria exceed 1, there is a relatively small chance of failure in the load layer and numerically zero chance of failure in the temperature layer.

In Figure 15, we plot the dominant sources of uncertainty in the nozzle predictions. Specifically, we identify the parameters and parameter combinations that have large Sobol indices. The parameter combinations depicted contribute a portion of the first 99.9% of variance of at least one nozzle QoI. It is interesting to note that, for each QoI, one parameter acting independently of all other parameters contributes over 90% of the total variance. Moreover almost all contributions to 99.9% of the variance of each QoI are caused by individual parameters. Only two second-order interactions are significant. However, the convergence curves in Figure 12 suggest that resolving multivariate interactions becomes important as we drive the surrogate error to tight tolerances. This is confirmed by further investigation of high-order Sobol indices that are small but nonzero. We omit these values for brevity.

The most sensitive parameter identified for each QoI matches physical intuition. The density of the CMC heat layer ρ^{CMC} is the largest contribution to the uncertainty in the mass of the nozzle, consistent with the observations that the minimum density of this layer is an order of magnitude larger than the maximum density of the other layers and the volume of this layer is greater than any other layer. The inlet stagnation pressure P_i has the greatest influence on thrust. The heat transfer coefficient β which transfers heat from the nozzle to the environment has the greatest impact on the




FIGURE 15 Sobol indices that contribute a portion of the first 99.9% of the variance of at least one nozzle QoI. The labels of the x-axis represent the parameter of the Sobol index. The maximum value of any Sobol index is one. We zoom in on the y-scale to better highlight the smaller Sobol index values, and report values greater than the 0.1 plot scale in the gray boxes. The definition of the parameter symbols can be found in Appendix A

Load layer temperature failure ratio, which is the ratio of the temperature in the load layer to the maximum allowable temperature. When β is decreased, less heat is able to escape the nozzle and thus the temperature in the nozzle increases. Finally, the failure strain in the CMC heat layer ϵ_f^{CMC} has the greatest impact on the failure of the thermal failure.

The sensitivity analysis presented here suggests that uncertainty in the model can be reduced by obtaining data that constrains uncertainty on three parameters, ρ^{CMC} , β and ϵ_f^{CMC} . The inlet stagnation temperature is a function of the flight conditions and so its variance cannot be constrained unless the aircraft mission is modified. Collecting experimental data and using Bayesian inference to infer updated parameter distributions to reduce uncertainty in model predictions is a topic for future work.

7 | CONCLUSIONS

We have described an integrated sparse grid approach for uncertainty quantification and sensitivity analysis, which generates a surrogate model that spans both random dimensions and model resolution dimensions. Novel aspects of this work include the extension from moment estimation to function approximation, extension to an adaptive multi-index construction (from MISC to AMISC), and demonstration of more complete UQ workflows for realistic applications, including computation of global sensitivity analysis and probability density functions for a jet engine nozzle application.

AMISC takes a possible set of fidelity levels and greedily allocates samples to the models which provide the most predictive utility per unit cost. The set of possible models must be specified before the algorithm is initiated. Given a vector of values for each discretization parameter of a model, we defined the set of possible models to be those indexed by the Cartesian product of these vectors. As an example, the advection diffusion model considered in this paper had three discretization parameters which controlled the accuracy and cost of the simulation. Each one of these hyper-parameters could take 6 values resulting in $6^3 = 216$ possible model fidelities.

Given a Cartesian product of possible model discretizations, AMISC applies adaptive generalized sparse grid refinement to select the tensor product approximation, from among multiple grid refinement candidates, that maximizes benefit per unit cost. Reflecting a focus on complex engineering applications with multiple statistics of interest, we used an adaptive criterion based upon a convex combination of mean and variance, across a vector of QoI, to guide refinement.

In the advection-diffusion model problem presented, AMISC was able to effectively allocate resources to the downward closed subset of the 216 model fidelities. The total computational cost required by AMISC to construct a surrogate with specified accuracy was over two orders of magnitude smaller relative to a sparse grid applied to only the most resolved model. Moreover, an instructive transition from early resolution of stochastic error to late resolution of deterministic

discretization bias is evident in the convergence history, demonstrating how multiple error sources are managed by the algorithm.

In the nozzle example, we show a variable level of improvement in efficiency for the four different QoI, where the efficiency is strongly dependent on the nonlinearity of the response. With greater nonlinearity, the AMISC approach displays strong savings. However, savings are not evident for nearly linear functions due to the additional overhead of carrying multiple models; rather, benefits are limited to greater flexibility in constructing rough approximations for low cost targets. Some inefficiencies were also caused by the downward-closed index set constraint for cases where final reductions in deterministic discretization error could only be addressed using the most resolved model, resulting in some loss in relative benefits for the tails of the convergence histories. Finally, we also demonstrate several downstream UQ products for the jet nozzle problem, including Sobol indices and probability density functions for each QoI, generated by processing and interrogating the final AMISC approximations.

In this paper, the relative cost of each model was determined using a small number of pilot runs prior to using AMISC. In future work, we intend to modify AMISC to determine relative model cost online during the evolution of the algorithm. Similarly, this paper focused on the case when the number of possible model discretizations is fixed a priori, as was necessary in the nozzle application. That is, we consider the highest-fidelity model to be most accurate and our aim is to efficiently construct a surrogate of the statistics of this model by including evaluations of a structured set of lower-fidelity models. In the future, we intend to extend our approach to consider an infinite set of possible discretizations, which converge to the exact solution to the governing equations. Finally, we intend to explore approaches that alter the downward-closed requirement, in order to allow additional refinement freedom. This is expected to be particularly helpful when model discrepancies do not demonstrate a smooth, monotonic decay.

ACKNOWLEDGMENTS

All authors were supported by the DARPA program for Enabling Quantification of Uncertainty in Physical Systems (EQUIPS). Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. The views expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

ORCID

John D. Jakeman  <https://orcid.org/0000-0002-3517-337X>

REFERENCES

1. Ghanem RG, Spanos PD. *Stochastic Finite Elements: A Spectral Approach*. New York, NY: Springer-Verlag; 1991.
2. Sudret B. Global sensitivity analysis using polynomial chaos expansions. *Reliab Eng Syst Safe*. 2008;93(7):964-979.
3. Xiu D, Karniadakis GE. The Wiener-Askey Polynomial Chaos for stochastic differential equations. *SIAM J Sci Comput*. 2002;24(2):619-644.
4. Arnst M, Ghanem R, Phipps E, Red-Horse J. Reduced chaos expansions with random coefficients in reduced-dimensional stochastic modeling of coupled problems. *Int J Numer Methods Eng*. 2014;97(5):352-376. <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.4595>.
5. Rasmussen CE, Williams CKI. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. Cambridge, MA: The MIT Press; 2005.
6. Sacks J, Welch WJ, Mitchell TJ, Wynn HP. Design and analysis of Computer Experiments. *Statist Sci*. 1989;4(4):409-423. <http://www.jstor.org/stable/2245858>.
7. Doostan A, Validi A, Iaccarino G. Non-intrusive low-rank separated approximation of high-dimensional stochastic models. *Comput Methods Appl Mech Eng*. 2013;263:42-55.
8. Gorodetsky A, Jakeman JD. Gradient-based optimization for regression in the functional tensor-train format. *J Comput Phys*. 2018;374:1219-1238. <http://www.sciencedirect.com/science/article/pii/S0021999118305321>.
9. Oseledets IV. Tensor-train decomposition. *SIAM J Sci Comput*. 2011;33(5):2295-2317.
10. Gorodetsky AA, Karaman S, Marzouk YM. Function-train: a continuous analogue of the tensor-train decomposition. *arXiv preprint arXiv:151009088*; 2015.
11. Nobile F, Tempone R, Webster CG. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J Num Anal*. 2008;46(5):2309-2345. <http://link.aip.org/link/?SNA/46/2309/1>.
12. Xiu D, Hesthaven JS. High-order collocation methods for differential equations with random inputs. *SIAM J Sci Comput*. 2005;27(3):1118-1139. <http://epubs.siam.org/doi/abs/10.1137/040615201>.
13. Jakeman JD, Roberts SG. Local and dimension adaptive stochastic collocation for uncertainty quantification. In: Garcke J, Griebel M, eds. *Sparse Grids and Applications, vol. 88 of Lecture Notes in Computational Science and Engineering*. Berlin, Germany: Springer; 2013:181-203 https://doi.org/10.1007/978-3-642-31703-3_9.

14. Agarwal N, Aluru NR. A data-driven stochastic collocation approach for uncertainty quantification in MEMS. *Int J Numer Methods Eng*. 2010;83(5):575-597. <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.2844>.
15. Cui T, Marzouk YM, Willcox KE. Data-driven model reduction for the Bayesian solution of inverse problems. *Int J Numer Methods Eng*. 2014;102(5):966-990. <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.4748>.
16. Soize C, Farhat C. A nonparametric probabilistic approach for quantifying uncertainties in low-dimensional and high-dimensional nonlinear models. *Int J Numer Methods Eng*. 2017;109(6):837-888. <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.5312>.
17. Carlberg K. Adaptive h-refinement for reduced-order models. *Int J Numer Methods Eng*. 2014;102(5):1192-1210. <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.4800>.
18. Perdikaris P, Raissi M, Damianou A, Lawrence ND, Karniadakis GE. Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proc R Soc Lond A*. 2017;473(2198):20160751. <http://rspa.royalsocietypublishing.org/content/473/2198/20160751>.
19. Kennedy MC, O'Hagan A. Predicting the output from a complex Computer code when fast approximations are available. *Biometrika*. 2000;87(1):1-13. <http://www.jstor.org/stable/2673557>.
20. Gratiet LL, Garnier J. Recursive co-kriging model for design of computer experiments with multiple levels of fidelity. *Int J Uncertain Quan*. 2014;4(5):365-386.
21. Narayan A, Gittelson C, Xiu D. A stochastic collocation algorithm with multifidelity models. *SIAM J Sci Comput*. 2014;36(2):A495-A521. <https://doi.org/10.1137/130929461>.
22. Teckentrup A, Jantsch P, Webster C, Gunzburger M. A multilevel stochastic collocation method for partial differential equations with random input data. *SIAM/ASA J Uncertain Quan*. 2015;3(1):1046-1074. <https://doi.org/10.1137/140969002>.
23. Haji-Ali A, Nobile F, Tamellini L, Tempone R. Multi-index stochastic collocation for random PDEs. *Comput Methods Appl Mech Eng*. 2016;306:95-122. <http://www.sciencedirect.com/science/article/pii/S0045782516301141>.
24. Gorodetsky A, Geraci G, Eldred M, Jakeman JD. A generalized framework for approximate control variates. *ArXiv*; 2018.
25. Peherstorfer B, Willcox K, Gunzburger M. Optimal model management for multifidelity Monte Carlo estimation. *SIAM J Sci Comput* *Appear*. 2016;38:A3163-A3194. <https://doi.org/10.1137/15M1046472>
26. Ng L, Eldred M. Multifidelity uncertainty quantification using non-intrusive polynomial Chaos and stochastic collocation. *Proceedings of 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference Honolulu, Hawaii*. Reston, VA: American Institute of Aeronautics and Astronautics; 2012.
27. Eldred MS, Ng LWT, Barone MF, Domino SP. In: Ghanem R, Higdon D, Owhadi H, eds. *Multifidelity Uncertainty Quantification Using Spectral Stochastic Discrepancy Models*. Cham: Springer International Publishing; 2016:1-45 https://doi.org/10.1007/978-3-319-11259-6_25-1.
28. Cliffe KA, Giles MB, Scheichl R, Teckentrup AL. Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Comput Visual Sci*. 2011;14(1):3 <https://doi.org/10.1007/s00791-011-0160-x>.
29. Geraci G, Eldred MS, Iaccarino G. A multifidelity multilevel Monte Carlo method for uncertainty propagation in aerospace applications. *19th AIAA Non-Deterministic Approaches Conference*. Reston, VA: American Institute of Aeronautics and Astronautics; 2017.
30. Giles MB. Multilevel Monte Carlo path simulation. *Oper Res*. 2008;56(3):607-617.
31. Haji-Ali AL, Nobile F, Tempone R. Multi-index Monte Carlo: when sparsity meets sampling. *Numer Math*. 2016;132(4):767-806. <https://doi.org/10.1007/s00211-015-0734-5>.
32. Ng L. Multifidelity approaches for Design under uncertainty. PhD thesis, aeronautics and astronautics, Massachusetts institute of. *Dent Tech*. 2013.
33. Fenrich R, Menier V, Avery P, Alonso J. Reliability-based design optimization of a supersonic nozzle. *Proceedings of the 6th European Conference on Computational Mechanics Glasgow*. Barcelona, Spain: International Center for Numerical Methods in Engineering; 2018.
34. Economon TD, Palacios F, Copeland SR, Lukaczyk TW, Alonso JJ. SU2: an open-source suite for multiphysics simulation and design. *AIAA J*. 2016;54(3):828-846.
35. Farhat C. *AERO-S: A General-Purpose Finite Element Structural Analyzer*; 2018. <https://bitbucket.org/frg/aero-s>.
36. Narayan A, Jakeman JD. Adaptive Leja sparse grid constructions for stochastic collocation and high-dimensional approximation. *SIAM J Sci Comput*. 2014;36(6):A2952-A2983. <https://doi.org/10.1137/140966368>.
37. Barthelmann V, Novak E, Ritter K. High dimensional polynomial interpolation on sparse grids. *Adv Comput Math*. 2000;12:273-288.
38. Bungartz HJ, Griebel M. Sparse grids. *Acta Numer*. 2004;13:147-269. <https://doi.org/10.1017/S0962492904000182>.
39. Smolyak SA. Quadrature and interpolation formulas for tensor products of certain classes of functions. *Soviet Math Dokl*. 1963;4:240-243.
40. Novak E, Ritter K. Simple cubature formulas with high polynomial exactness. *Const Approx*. 1999;15(4):499-522. <https://doi.org/10.1007/s003659900119>.
41. de Baar JHS, Roberts SG. Multifidelity sparse-grid-based uncertainty quantification for the Hokkaido Nansei-oki tsunami. *Pure Appl Geophys*. 2017;174(8):3107-3121. <https://doi.org/10.1007/s00024-017-1606-y>.
42. Gerstner T, Griebel M. Dimension-adaptive tensor-product quadrature. *Comput Secur*. 2003;71(1):65-87.
43. Farcaş IG, Sărbu PC, Bungartz HJ, Neckel T, Uekermann B. Multilevel adaptive stochastic collocation with dimensionality reduction. In: Garcke J, Pflüger D, Webster CG, Zhang G, eds. *Sparse Grids and Applications – Miami 2016*. Cham: Springer International Publishing; 2018:43-68.
44. Sobol' IM. Sensitivity estimates for nonlinear mathematical models. *Math Model Comput Exp*. 1993;1(4):407-414.
45. Saltelli A, Chan K, Scott EM. *Sensitivity Analysis*. New York: Wiley; 2000.
46. Buzzard GT. Global sensitivity analysis using sparse grid interpolation and polynomial chaos. *Reliab Eng Syst Safe* 2012;107(0):82-89. <http://www.sciencedirect.com/science/article/pii/S09518320110%01530>.

47. Formaggia L, Guadagnini A, Imperiali I, et al. Global sensitivity analysis through polynomial chaos expansion of a basin-scale geochemical compaction model. *Comput Geosci*. 2013;17(1):25-42. <https://doi.org/10.1007/s10596-012-9311-5>.
48. Buzzard GT. Efficient basis change for sparse-grid interpolating polynomials with application to T-cell sensitivity analysis. *Comput Biol J*. 2013;2013: <http://www.hindawi.com/journals/cbj/2013/562767/abs/>.
49. Jakeman JD, Franzelin F, Narayan A, Eldred M, Plfüger D. Polynomial chaos expansions for dependent random variables. *Comput Methods Appl Mech Eng*. 2019;351:643-666. <http://www.sciencedirect.com/science/article/pii/S00457825193%01884>.
50. Butler T, Jakeman JD, Wildey T. Convergence of probability densities using approximate models for forward and inverse problems in uncertainty quantification. *ArXiv:180700375*; 2018.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Jakeman JD, Eldred MS, Geraci G, Gorodetsky A. Adaptive multi-index collocation for uncertainty quantification and sensitivity analysis. *Int J Numer Methods Eng*. 2020;121:1314–1343. <https://doi.org/10.1002/nme.6268>

APPENDIX NOZZLE MODEL INPUT UNCERTAINTIES

The materials of the nozzle, inlet conditions, atmospheric conditions, and heat transfer are all sources of uncertainty that effect predictions of our nozzle model. In this section we explicitly identify and characterize (assign a probability distribution) to each source of uncertainty.

Material constants

Figure A1 shows the individual wall layers in the nozzle wall. The material properties of each of these layers affect nozzle performance. The thermal insulation “aerogel” is not modeled, but all other layers are modeled. An incomplete set of material properties is gathered from a variety of experimental data sources.

Missing macroscopic properties are either determined through fundamental physical analysis of the material's structure or are approximated by values for similar materials. Table A1 provides the properties used in the MULTI - F analysis for the CMC material, Table A2 for the GR-BMI material, and Table A3 for the TI-HC material. Last, Table A4 lists the thermal conductivity of the air gap and gives the distribution for the assumed yield stress of the composite panel material used for the baffles.




FIGURE 1 Individual layers in the nozzle wall [Color figure can be viewed at wileyonlinelibrary.com]

TABLE A1 Assumed isotropic material properties for the heat layer's ceramic matrix composite material (CMC)

| Property | symbol | Units | Distribution |
|-------------------------------|-------------------------------|--------------------------------|-----------------------------------|
| Density | ρ^{CMC} | $\frac{\text{kg}}{\text{m}^3}$ | $\mathcal{U}(2293.789, 2496.486)$ |
| Elastic modulus | E^{CMC} | GPa | $\mathcal{U}(58.940, 76.163)$ |
| Poisson ratio | ν^{CMC} | | $\mathcal{U}(0.23, 0.43)$ |
| Thermal conductivity | k^{CMC} | $\frac{\text{W}}{\text{m-K}}$ | $\mathcal{U}(1.37, 1.45)$ |
| Thermal expansion coefficient | α^{CMC} | $\text{K}^{-1} \times 10^{-6}$ | $\mathcal{U}(0.228, 0.252)$ |
| Max service temperature | $T_{\text{max}}^{\text{CMC}}$ | K | $\mathcal{U}(963, 983)$ |
| Failure strain | ϵ_f^{CMC} | % | $\mathcal{U}(0.049788, 0.096441)$ |

| Property | Symbol | Units | Distribution |
|-------------------------------------|---------------------------------------|--------------------------------|---------------------|
| Density | ρ^{GR} | $\frac{\text{kg}}{\text{m}^3}$ | $U(1563, 1573)$ |
| Elastic moduli | $E_1 = E_2$ | GPa | $U(57, 63)$ |
| In-plane shear modulus | G_{12} | GPa | $U(22.6, 24.0)$ |
| Poisson ratios | $\nu_{12} = \nu_{21}$ | | $U(0.334, 0.354)$ |
| Mutual influence coef (first kind) | $\mu_{1,12} = \mu_{2,12}$ | | $U(-0.1, 0.1)$ |
| Mutual influence coef (second kind) | $\mu_{12,1} = \mu_{12,2}$ | | $U(-0.1, 0.1)$ |
| Thermal conductivity | $k_1 = k_2$ | $\frac{\text{W}}{\text{m-K}}$ | $U(3.208, 3.546)$ |
| Thermal conductivity | k_3 | $\frac{\text{W}}{\text{m-K}}$ | $U(3.243, 3.585)$ |
| Thermal expansion coef | $\alpha_1 = \alpha_2$ | $\text{K}^{-1} \times 10^{-6}$ | $U(1.16, 1.24)$ |
| Thermal expansion coef | α_{12} | $\text{K}^{-1} \times 10^{-6}$ | $U(-0.04, 0.04)$ |
| Max service temperature | $T_{\text{max}}^{\text{GR}}$ | K | $U(500, 510)$ |
| Failure strain (tension) | $\epsilon_{f,1}^t = \epsilon_{f,2}^t$ | % | $U(0.675, 0.825)$ |
| Failure strain (compression) | $\epsilon_{f,1}^c = \epsilon_{f,2}^c$ | % | $U(-0.572, -0.494)$ |
| Failure strain (shear) | γ_f | % | $U(0.153, 0.187)$ |

TABLE A2 Macroscopic laminate material properties for the load layer's graphite/bismaleimide (GR-BMI) composite layers used in the axisymmetric shell specification

| Property | Symbol | Units | Distribution |
|-------------------------------|------------------------------|--------------------------------|---------------------|
| Density | ρ^{TI} | $\frac{\text{kg}}{\text{m}^3}$ | $U(177.77, 181.37)$ |
| Elastic modulus | E^{TI} | GPa | $U(1.587, 2.823)$ |
| Poisson ratio | ν^{TI} | | $U(0.160, 0.196)$ |
| Thermal conductivity | k^{TI} | $\frac{\text{W}}{\text{m-K}}$ | $U(0.680, 0.736)$ |
| Thermal expansion coefficient | α^{TI} | $\text{K}^{-1} \times 10^{-6}$ | $U(2.88, 3.06)$ |
| Max service temperature | $T_{\text{max}}^{\text{TI}}$ | K | $U(745, 765)$ |
| Yield stress | σ_Y^{TI} | MPa | $U(9.676, 16.951)$ |

TABLE A3 Assumed isotropic macroscopic material properties for titanium honeycomb layer (TI-HC)

| Property | Symbol | Units | Distribution |
|----------------------|------------------|-------------------------------|---------------------|
| Thermal conductivity | k^{gap} | $\frac{\text{W}}{\text{m-K}}$ | $U(0.0320, 0.0530)$ |
| Yield strength | $\sigma_{Y,B}$ | MPa | $U(56.876, 99.219)$ |

TABLE A4 Assumed properties of air gap between thermal and load layers and of the panel structure used in baffles

TABLE A5 Mission parameters

| Parameter | Symbol | Units | Nominal Value | Distribution |
|--|---------|--|----------------------------|--------------|
| Inlet stagnation pressure | P_i | Pa | $U(86362.122, 113062.433)$ | |
| Inlet stagnation temperature | T_i | K | $U(928.729, 981.600)$ | |
| Atmospheric pressure | P_a | Pa | $U(17386.668, 20206.125)$ | |
| Atmospheric temperature | T_a | K | $U(202.856, 231.297)$ | |
| Heat transfer coefficient to environment | β | $\frac{\text{W}}{\text{m}^2\text{-K}}$ | $U(4.043, 37.376)$ | |

Mission parameters

We focus on a typical reconnaissance mission for a small high-subsonic unmanned military aircraft. The mission includes climbing at maximum rate to a cruise altitude of 43 000 ft, cruising at Mach 0.92 for a specified distance to an observation point (eg, 500 km), loitering at an altitude of 43 000 ft and Mach 0.5 for 2 hours, and then returning to the takeoff point. On the return, the aircraft descends to 10 000 ft and cruises at Mach 0.9 in a high-speed “dash” segment lasting several kilometers before landing. The analysis of this mission showed that the climb segment was the most critical for nozzle performance since maximum thrust was required at all altitudes, leading to the highest temperatures and pressures at the inlet of the nozzle. In particular, the state of climb right before beginning the cruise segment was the most critical in terms of stresses and temperatures experienced by the nozzle. Consequently, we focus on the climb segment of the mission here. Table A5 summarizes the mission parameters used in the nozzle analysis. We assume the altitude and mach number are fixed at 40 000 ft and 0.511, respectively.