


# Noise-robust multi-fidelity surrogate modelling for parametric partial differential equations

Benjamin M. Kent <sup>a,\*</sup>, Lorenzo Tamellini <sup>a</sup>, Matteo Giacomini <sup>b,c</sup>,  
Antonio Huerta <sup>b,c</sup>

<sup>a</sup> Consiglio Nazionale delle Ricerche – Istituto di Matematica Applicata e Tecnologie Informatiche “E. Magenes” (CNR-IMATI), Via Adolfo Ferrata, 5, Pavia, PV, 27100, Italy

<sup>b</sup> Laboratori de Càlcul Numèric (LaCàN), ETS de Ingeniería de Caminos, Canales y Puertos, Universitat Politècnica de Catalunya - BarcelonaTech (UPC), Barcelona, Spain

<sup>c</sup> Centre Internacional de Mètodes Numèrics en Enginyeria (CIMNE), Barcelona, Spain

## ARTICLE INFO

### Keywords:

Surrogate modelling  
Multi-index stochastic collocation  
Noise robust  
Noisy solvers  
Uncertainty quantification

## ABSTRACT

We address the challenge of constructing noise-robust surrogate models for quantities of interest (QoIs) arising from parametric partial differential equations (PDEs), using multi-fidelity collocation techniques; specifically, the Multi-Index Stochastic Collocation (MISC). In practical scenarios, the PDE evaluations used to build a response surface are often corrupted by numerical noise, especially for the low-fidelity models. This noise, which may originate from loose solver tolerances, coarse discretisations, or transient effects, can lead to overfitting in MISC, degrading surrogate quality through nonphysical oscillations and loss of convergence, thereby limiting its utility in downstream tasks like uncertainty quantification, optimisation, and control. To correct this behaviour, we propose an improved version of MISC that can automatically detect the presence of solver noise during the surrogate model construction and then ignore the exhausted fidelities. Our approach monitors the spectral decay of the surrogate at each iteration, identifying stagnation in the coefficient spectrum that signals the onset of noise. Once detected, the algorithm selectively halts the use of noisy fidelities, focusing computational resources on those fidelities that still provide meaningful information. The effectiveness of this approach is numerically validated on two challenging test cases: a parabolic advection–diffusion PDE with uncertain coefficients, and a parametric turbulent incompressible Navier–Stokes problem. The results showcase the accuracy and robustness of the resulting multi-fidelity surrogate and its capability to extract relevant information, even from under-resolved meshes not suitable for reliable single-fidelity computations.

## 1. Introduction

Parametric partial differential equations (PDEs) are commonly used to model physical problems with user-defined, possibly uncertain, configurations. The dependence of the solution on such (uncertain) parameters, or that of a quantity of interest derived from the solution, is often well approximated using global polynomial approximation under mild assumptions on the parameter dependence of the problem. Popular constructions include using sparse grids, radial basis functions, Gaussian processes, least squares methods, stochastic Galerkin, reduced basis methods, proper orthogonal and proper generalized decompositions and neural networks [1–13].

\* Corresponding author.

E-mail addresses: [kent@imati.cnr.it](mailto:kent@imati.cnr.it), [benkent@live.co.uk](mailto:benkent@live.co.uk) (B.M. Kent), [tamellini@imati.cnr.it](mailto:tamellini@imati.cnr.it) (L. Tamellini).

Many of these constructions require repeatedly solving the PDE. To reduce computational costs, so-called multi-fidelity variants of these methods have been introduced. These may use greedy adaptive algorithms in both parametric and physical (space, time...) domains [14–17], or multi-mesh and multi-level approaches [18–23]; see e.g., [24] for a survey.

This work concerns multi-fidelity approximation. Besides the already-hinted mesh or timestep refinement, a multi-fidelity approach could also consider solvers with different tolerances, or even a set of different physical models of the system at hand. Specifically, we focus on *multi-index stochastic collocation* (MISC) as introduced in [25]. This is a multi-fidelity surrogate modelling technique that builds upon the well-known sparse grid stochastic collocation (SGSC) approximation method for parametric PDEs [26–29]. Methodologies built upon SGSC are popular due to their non-intrusive nature: a quantity of interest needs to be evaluated only for a prescribed collocation set of parameter realisations, and the surrogate model is then built by interpolating these evaluations. In our multi-fidelity context, the collocation set specifies both parameter realisations and solver fidelities, which are combined in a similar way to SGSC methods to create a multi-fidelity surrogate model. Adaptive MISC techniques have been successfully applied in a number of practical applications, see e.g., [18,30,31].

Adaptive SGSC and MISC algorithms generally use a “greedy” approach that refines the surrogate model by seeking the greatest change to the computed surrogate. This assumes that every change in the surrogate is a feature that must be captured to achieve an accurate final approximation. In many applications this is however not the case: the underlying solvers (especially the lowest fidelities) may introduce spurious, uncontrolled errors (e.g., due to iterative solver error, discretisation errors,...). Hence greedy strategies may be unable to distinguish between the *true* solution behaviour and the solver error, and consequently “overfit” to the inherent solver noise. This leads to spurious oscillations in the resulting surrogate models. See for example [30,32–34], as well as [35,36] and references therein for further discussion on noise corruption in the context of outer loop applications such as uncertainty quantification (UQ) or shape optimisation.

Classical least-squares approaches (either single- or multi-fidelity) to surrogate modelling construction can mitigate the impact of noisy evaluations of the quantity of interest [37–42], although this noise is not always explicitly accounted for in these approaches; the impact of noise in the convergence of least squares approximations is discussed in e.g. [43–46]. Among the alternative approaches for dealing with noise in model evaluations, we mention stochastic radial basis functions [30,34], the approximate control variate method for non-deterministic models discussed in [36], and the noise-robust approximation of level sets of multivariate functions discussed in [47].

With the same purpose, in this work we propose instead a novel noise-robust variant of the MISC algorithm, that we call PLATEAUMISC. PLATEAUMISC modifies the greedy MISC algorithm to differentiate between refinements that capture the *true* parametric response and refinements that are capturing solver noise. We achieve this by inspecting the spectral content of the approximation and automatically detecting deviations of the behavior of the multi-fidelity response surfaces from the expected spectral decay of smooth quantities of interest, see e.g. [7, Chapter 4]. More precisely, the algorithm detects the corruption due to solver error by a failure to adhere to an a priori bound for the spectral polynomial expansion coefficients.

Compared to the classical version of MISC, the PLATEAUMISC algorithm thus offers a multi-fidelity approximation with the following advantages:

- it is robust to solver noise, allowing the exploitation of information from low fidelity approximations without corrupting the final approximation;
- it automatically and efficiently transitions from low-fidelity to high-fidelity solvers, blocking the “exhausted” low fidelity solvers, as the computational budget increases.


We conclude this introduction by briefly mentioning a related line of research on fault-tolerant surrogate model construction. This approach addresses settings in which a *small* number of evaluations of the quantity of interest are severely *corrupted* or even *unavailable* due to simulation failures. This contrasts with the scenario considered in the present work, where we assume that a *large* number of evaluations are affected by *noise*. Fault-tolerant surrogate modeling techniques based on compressed sensing are discussed in [48,49] and could, in principle, be adapted to the setting studied here, whereas the fault-tolerant sparse grid approximations proposed in [50,51] do not extend to our scenario; see also [52] for a fault-tolerant multi-level Monte Carlo method.

The paper is structured as follows. Section 2 details the construction of a single-fidelity SGSC approximation and the polynomial approximation properties that will be exploited. Section 3 presents the classic MISC algorithm and the novel PLATEAUMISC adaptive algorithm exploiting spectral coefficient plateaus. Section 4 numerically demonstrates the algorithm applied to a parabolic advection–diffusion PDE with uncertain parameters and a parametric Reynolds Averaged Navier–Stokes flow with Spalart–Allmaras turbulence model. The paper concludes in Section 5.

## 2. Collocation based methods for parametric PDEs

In abstract terms, the purpose of this work is to accurately and efficiently approximate a multi-variate function  $q : \Gamma \subset \mathbb{R}^d \rightarrow \mathbb{R}$  whose values are only known implicitly upon solving a PDE. The inputs  $\mathbf{y} \in \Gamma$  will be referred to as parameters and we consider a Cartesian parameter domain  $\Gamma := \Gamma_1 \times \Gamma_2 \times \dots \times \Gamma_d$ . For simplicity it is assumed  $\Gamma_i = [0, 1]$  for  $i = 1, 2, \dots, d$  and that the quantity of interest (QoI)  $q$  is a scalar-valued function. For instance, in the framework of uncertainty quantification the parameters represent images of underlying random variables (Section 4.1), whereas in parametric flow problems they describe working conditions of the system (Section 4.2).

Assuming that our quantity of interest is smooth with respect to the parameters  $\mathbf{y}$ , global polynomial approximation represents a common choice in the context of parametric PDEs and related problems, as already mentioned in the introduction. In particular, this



**Fig. 1.** Example sequences of one-dimensional collocation points  $Z^\beta$  on  $[0, 1]$  for Clenshaw–Curtis points with the ‘doubling’ rule and symmetric Leja points with the ‘two-step’ rule, as defined in (2).

assumption ensures that we can write the quantity of interest as a convergent polynomial series. More explicitly, consider a basis of multi-variate polynomials  $\{\Phi_{\underline{p}}(\mathbf{y})\}_{\underline{p} \in \mathbb{N}_0^d}$  where  $\Phi_{\underline{p}}(\mathbf{y}) = \prod_{i=1}^d \phi_{p_i}(y_i)$  and each term  $\phi_{p_i}(y_i)$  is a polynomial of degree  $p_i$  in the parameter  $y_i$ . For example,  $\phi_{p_i}(y_i)$  could be the monomial  $y_i^{p_i}$ , or the Chebyshev polynomial of the first kind of degree  $p_i$ . The quantity of interest  $q$  can then be expressed as a convergent polynomial series,

$$q(\mathbf{y}) \equiv \sum_{\underline{p} \in \mathbb{N}_0^d} \hat{q}_{\underline{p}} \Phi_{\underline{p}}(\mathbf{y}). \tag{1}$$

In this work we consider global polynomial approximations constructed via sparse grid stochastic collocation (SGSC) type methods which we now briefly recall, see e.g. [27,29,53]. In particular, we closely follow the setup of [29].

*One dimensional interpolation.* First, denote by  $\{Z^\beta\}_{\beta \in \mathbb{N}}$  a sequence of sets of one dimensional collocation points where  $Z^\beta := \{z_i\}_{i=1}^{m_\beta} \subset [0, 1]$ . The number of points  $m_\beta$  at each level  $\beta \in \mathbb{N}$  is specified by the so-called *level-to-knots* function. Common choices of the level-to-knots function  $m_\beta$  are

$$m_\beta = 1 \text{ for } \beta = 1, \text{ and } m_\beta = \begin{cases} \beta & \text{‘linear’ rule} \\ 2(\beta - 1) + 1 & \text{‘two-step’ rule for all } \beta > 1. \\ 2^{\beta-1} + 1 & \text{‘doubling’ rule} \end{cases} \tag{2}$$

A common choice of points  $Z^\beta := \{z_i\}_{i=1}^{m_\beta}$  are the Clenshaw–Curtis (CC) points which are defined by the extrema of the Chebyshev polynomials [54]. CC points are nested under the ‘doubling’ level-to-knots rule of (2). In this work, we also use the symmetric Leja points for a uniform random variable. By construction, Leja points are a nested sequence of points for every choice of level-to-knots rule in (2) [29]. Clenshaw–Curtis points, symmetric Leja points and the ‘doubling’ and ‘two-step’ rules are illustrated in Fig. 1.

Next, we define a corresponding sequence of interpolation operators  $\{I^\beta\}_{\beta \in \mathbb{N}}$  such that for continuous  $q \in C^0([0, 1]; \mathbb{R})$  it holds that

$$I^\beta[q](y) := \sum_{z \in Z^\beta} q(z) L_z^{Z^\beta}(y), \tag{3}$$

where  $\{L_z^{Z^\beta}(y)\}_{z \in Z^\beta}$  are the corresponding Lagrange interpolation polynomials  $L_z^{Z^\beta}(y) = \prod_{w \in Z^\beta \setminus \{z\}} (y - w)/(z - w)$ .

*Multi-dimensional interpolation.* A multi-dimensional approximation is constructed using the multi-index set  $I \subset \mathbb{N}^d$  to define a linear combination of tensor products of interpolation operators. This is known to be an effective approach for approximating suitably smooth high-dimensional functions ( $d \gg 1$ ) without resorting to tensor grid sampling of the parameter domain  $\Gamma$  which quickly becomes computationally unaffordable (i.e. *the curse of dimensionality*) [55]. An approximation of a multi-variate function  $q \in C^0(\Gamma; \mathbb{R})$  is defined via the *combination technique* formulation [56,57],

$$q(\mathbf{y}) \approx q^I(\mathbf{y}) := I^I[q](\mathbf{y}) := \sum_{\underline{\beta} \in I} c_{\underline{\beta}} I^{\underline{\beta}}[q](\mathbf{y}), \tag{4}$$

where the combination technique coefficients  $c_{\underline{\beta}}$  and tensorisation of the interpolation operators (3) are defined by

$$c_{\underline{\beta}} := \sum_{\substack{\mathbf{v} \in \{0,1\}^d \\ \underline{\beta} + \mathbf{v} \in I}} (-1)^{\|\mathbf{v}\|_1}, \quad I^{\underline{\beta}}[q](\mathbf{y}) := \bigotimes_{i=1}^d I^{\beta_i}[q](y_i). \tag{5}$$

Eq. (4) provides a global polynomial approximation over the parameter domain  $\Gamma$ , within the polynomial approximation space  $\mathcal{P}^I(\Gamma; \mathbb{R})$  determined by the multi-index set  $I$ . Denote the set of  $m_I$  distinct collocation points

$$Z^I := \bigcup_{\substack{\underline{\beta} \in I \\ c_{\underline{\beta}} \neq 0}} Z^{\beta_1} \times Z^{\beta_2} \times \dots \times Z^{\beta_d}. \tag{6}$$

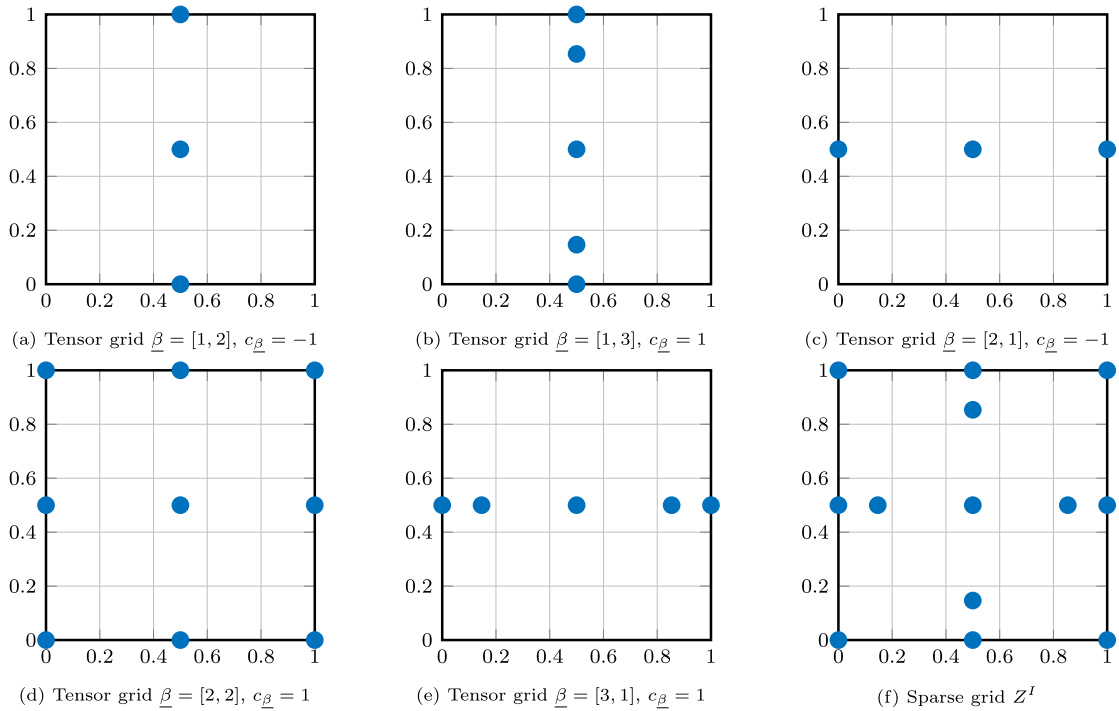


Fig. 2. For a multi-index set  $I = \{\underline{\beta} \in \mathbb{N}^2 \text{ s.t. } \|\underline{\beta}\|_1 \leq 2 + 2\}$  using Clenshaw–Curtis points with the doubling rule from (2), the tensor grids with non-zero combination technique coefficient  $c_{\underline{\beta}}$  and the complete sparse grid  $Z^I$  defined via (6) are shown in Fig. 2a–e and f respectively.




Fig. 3. Example reduced margin  $R_I$  for a multi-index set with  $d = 2$ .

Note that if the sequence of points  $\{Z^\beta\}_{\beta \in \mathbb{N}}$  is nested, i.e.  $Z^1 \subset Z^2 \subset Z^3 \subset \dots$ , then many points in (6) are coincident. The sparse grid construction (6) is demonstrated with a  $d = 2$  dimensional example in Fig. 2.

**Admissibility of multi-index sets.** A valid approximation in (4) is obtained only if the multi-index set  $I$  is admissible (also known as downwards closed, or a lower set). Admissibility requires each multi-index  $\underline{\beta} \in I$  to satisfy

$$\underline{\beta} - \underline{e}_i \in I \text{ for all } i = 1, \dots, d \text{ s.t. } \beta_i > 1 \tag{7}$$

where  $\underline{e}_i = [0, 0, \dots, 1, \dots, 0]$  is the unit multi-index with value 1 only in the  $i^{\text{th}}$  entry. The set of multi-indices that can be added to  $I$  whilst maintaining admissibility is called the reduced margin

$$R_I := \left\{ \underline{\beta} \in \mathbb{N}^d \setminus I \text{ s.t. } \underline{\beta} - \underline{e}_i \in I \text{ for all } i = 1, \dots, d \text{ s.t. } \beta_i > 1 \right\}. \tag{8}$$

For an admissible multi-index set  $I \subset \mathbb{N}^d$  and nested set of points  $\{Z^\beta\}_{\beta \in \mathbb{N}}$ , the approximation (4) is an interpolant. An example multi-index set  $I$  and the corresponding reduced margin  $R_I$  are shown in Fig. 3.

**Polynomial approximation spaces.** Recall that  $\mathcal{P}^I(\Gamma; \mathbb{R})$  denotes the polynomial space to which the approximation  $q^I$  in (4) belongs. This space can be obtained as the span of a suitable set of multi-variate polynomials, for example the multi-variate monomials, or multi-variate spectral-type polynomials, e.g. Legendre or Chebyshev polynomials. We require a spectral-type polynomial basis, and again we

denote by  $\{\phi_p\}_{p \in \mathbb{N}}$  a sequence of univariate polynomials of degree  $p$ , and let  $\Phi_{\underline{p}}(\mathbf{y}) = \prod_{i=1}^d \phi_{p_i}(y_i)$ . For the polynomial approximation space  $\mathcal{P}^I(\Gamma; \mathbb{R})$  to which the sparse grid approximation (4) belongs, defined implicitly for multi-index set  $I$ , a corresponding multi-index set of polynomial degrees  $G(I)$  can be defined such that

$$\mathcal{P}^I(\Gamma; \mathbb{R}) = \text{span}_{p \in G(I)} \{\Phi_{\underline{p}}\}, \tag{9}$$

i.e.  $G(I)$  is the appropriate multi-index set representing the multi-variate polynomial degrees. For an approximation  $q^I(\mathbf{y})$  computed via the combination technique, it is straightforward to perform a change of basis to the multi-variate polynomial basis  $\{\Phi_{\underline{p}}\}_{p \in G(I)}$  via the solution of a linear system [58,59]. This gives a *truncated* spectral polynomial expansion

$$q^I(\mathbf{y}) \equiv \sum_{\underline{p} \in G(I)} \hat{q}_{\underline{p}}^I \Phi_{\underline{p}}(\mathbf{y}). \tag{10}$$

The equivalence of the sparse grid formulation and the truncated spectral expansion is at the core of the PLATEAUMISC algorithm introduced in Section 3. Indeed, once the MISC approximation is constructed, the change of basis to spectral polynomials in (10) allows exploitation of the spectral polynomial coefficient decay properties in the constructed approximation. This information is thus employed to detect solver errors via the corruption of spectral polynomial coefficients (Section 2.1) and consequently devise a strategy to construct an approximant robust to noisy evaluations (Section 3).

### 2.1. Effect of solver errors in sparse grid stochastic collocation

It is generally not possible to evaluate  $q(\mathbf{y})$  exactly: each evaluation used to build (4) is an approximation  $q^\alpha(\mathbf{z}) \approx q(\mathbf{z})$  via a PDE solve where  $\alpha$  is a multi-index of hyperparameters that control the approximation fidelity (e.g. mesh size, solver tolerances,...). Therefore, a computable single-fidelity surrogate of the quantity of interest can be formed using the combination technique (4) with multi-index set  $I$ ,

$$q(\mathbf{y}) \approx q^{\alpha, I}(\mathbf{y}) := \sum_{\underline{\beta} \in I} c_{\underline{\beta}} \mathcal{I}^{\underline{\beta}}[q^\alpha](\mathbf{y}). \tag{11}$$

constructed using pointwise evaluations  $q^\alpha(\mathbf{z}) \approx q(\mathbf{z})$  for all  $\mathbf{z} \in Z^I$ . Again, as in (10), a truncated spectral expansion can be obtained via a change of basis,

$$q^{\alpha, I}(\mathbf{y}) \equiv \sum_{\underline{p} \in G(I)} \hat{q}_{\underline{p}}^{\alpha, I} \Phi_{\underline{p}}(\mathbf{y}). \tag{12}$$

Each computable coefficient  $\hat{q}_{\underline{p}}^{\alpha, I}$  for  $\underline{p} \in G(I)$  is an approximation of the corresponding coefficient  $\hat{q}_{\underline{p}}$  from the full spectral expansion (1) and can be expressed as

$$\hat{q}_{\underline{p}}^{\alpha, I} = \underbrace{\hat{q}_{\underline{p}}}_{(I)} + \underbrace{\hat{q}_{\underline{p}}^I - \hat{q}_{\underline{p}}}_{(II)} + \underbrace{\hat{q}_{\underline{p}}^{\alpha, I} - \hat{q}_{\underline{p}}^I}_{(III)}. \tag{13}$$

The term (I) in (13) is the *true* spectral series coefficient, encapsulating information on the smoothness of the quantity of interest  $q$ . It follows that the decay of (I) represents the theoretical behaviour of the spectral coefficients according to the regularity of  $q(\mathbf{y})$ . In particular, for an analytic function  $q(\mathbf{y})$  in a spectral-type polynomial expansion, we expect the expansion coefficients  $\{\hat{q}_{\underline{p}}\}_{\mathbb{N}_0^d}$  to obey an a priori bound of the form

$$|\hat{q}_{\underline{p}}| \leq C(q, \underline{p}) \prod_{i=1}^d e^{-g_i p_i} \tag{14}$$

where  $g_i > 0$  for all  $i = 1, 2, \dots, d$  and  $C(q, \underline{p})$  grows slowly as the polynomial degree  $\underline{p}$  increases. For further details on analyticity and spectral polynomial expansion coefficient bounds see, e.g. [60] and [7, Chapter 3, Theorem 3.2]. When the term (I) dominates (13) the coefficient is well approximated.

The term (II) represents an approximation error due to computing the spectral coefficient  $\hat{q}_{\underline{p}}^I$  using the polynomial interpolant (4) (rewritten in the form (10)), instead of the true coefficient  $\hat{q}_{\underline{p}}$  in the infinite expansion (1). This contains aliasing type effects due to using the subspace  $\mathcal{P}^I(\Gamma; \mathbb{R})$  in (9) implicitly defined by the interpolation. This term is at its largest for the highest polynomial degree terms in the approximation. The effect of (II) can be mitigated for a given polynomial degree  $p$  by increasing the size of the polynomial approximation space. These effects are generally not a source of concern in the context of this work.

The term (III) represents the solver error, which is usually considered as an intrinsic source of noise in simulation-based data. The proposed PLATEAUMISC strategy aims to robustify the standard MISC algorithm to handle this uncertainty. By linearity of the change of basis to the spectral polynomials and linearity of the combination technique construction (4) it follows that (III) is in fact the spectral coefficient of an interpolant of the solver error,

$$\sum_{\underline{p} \in G(I)} \left( \hat{q}_{\underline{p}}^{\alpha, I} - \hat{q}_{\underline{p}}^I \right) \Phi_{\underline{p}}(\mathbf{y}) \equiv q^{\alpha, I}(\mathbf{y}) - q^I(\mathbf{y}) \equiv \sum_{\underline{\beta} \in I} c_{\underline{\beta}} \mathcal{I}^{\underline{\beta}}[q^\alpha - q](\mathbf{y}). \tag{15}$$

If the solver error is constant over  $\Gamma$ , i.e.  $E^\alpha(\mathbf{y}) = (q^\alpha - q)(\mathbf{y}) = C^\alpha \in \mathbb{R}$ , then the interpolant on the right-hand side of (15) is exact and a constant. Therefore the only non-zero coefficient in the expansion on the left-hand side of (15) is for the constant polynomial, that is, the term (III) is non-zero only for  $p = 0$ . If instead the solver error is more complex, possibly even seemingly random over  $\Gamma$ , many coefficients  $\hat{q}_p^{\alpha,I} - \hat{q}_p^I$  will be non-zero and potentially large; for instance, this can be due to unresolved or preasymptotic meshes, coarse tolerances in iterative solvers, timestepping etc. It is this term (III) that can corrupt the computation of the surrogate  $q^{\alpha,I}$  in (11). This can result in a *spectral plateau*, i.e. a situation in which many coefficients corresponding to sufficiently high polynomial degree  $\|p\|_1$  terms in the expansion (10) are corrupted and of a similar erroneous magnitude. The following Example 1 provides further intuition for the expected stagnation in the decay of the spectral coefficients under standard assumptions on the quantity of interest.

**Example 1** (Estimating coefficient plateau for Gaussian noise). Suppose that the solver error  $E^\alpha(\mathbf{y})$  is a white Gaussian random field over  $\Gamma$ , i.e.,  $E^\alpha(\mathbf{y}) \sim \mathcal{N}(0, \sigma^2)$  and  $E^\alpha(\mathbf{y}_i)$  is statistically independent of  $E^\alpha(\mathbf{y}_j)$  for all  $\mathbf{y}_i, \mathbf{y}_j \in \Gamma$ . Consider the ‘linear’ rule in (2),  $m_\beta = \beta$ , and take the abscissae of the mid-point quadrature rule in  $[0, 1]$  to be the set of one-dimensional collocation points  $Z^\beta = \{z_i\}_{i=1}^\beta$ , i.e.,  $z_i = \frac{1}{\beta}(i - \frac{1}{2})$ ; finally, let  $I = [w, w, w, \dots]$  for some  $w \in \mathbb{N}, w \geq 1$ . The sparse grids interpolant reduces to a tensor-grid interpolant over a Cartesian grid  $Z^I = Z^w \times Z^w \times \dots = \{z_i\}_{i=1}^M$  with  $M = w^d$  collocation points, and the coefficients  $c_\beta$  in (4), in (11) and in the right-hand side of (15) are all zero other than  $c_{[w,w,\dots]} = 1$ . Furthermore, let  $\{\Phi_{p_i}\}_{i=1}^M$  be the Legendre orthonormal polynomials with total polynomial degree up to  $w - 1$ , i.e.  $G(I) = \{p : \|p\|_\infty \leq w - 1\}$  in the left-hand side of (15). We introduce a suitable ordering of  $G(I)$  such that we can index its elements as  $p_i$  and hence the set of polynomials is  $\{\Phi_{p_i}\}_{i=1}^M$  with a scalar index  $i$ . We can collect the pointwise errors  $\{E^\alpha(z_i)\}_{i=1}^M$  in a (random) vector

$$e := [E^\alpha(z_1), E^\alpha(z_2), \dots, E^\alpha(z_M)]^\top \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbb{1}), \tag{16}$$

where  $\mathbb{1}$  is the identity matrix in  $\mathbb{R}^{M \times M}$ . Similarly, we can collect the coefficients of the spectral expansion at the left-hand side of (15) in another vector

$$\hat{e} := [\hat{q}_{p_1}^{\alpha,I} - \hat{q}_{p_1}^I, \hat{q}_{p_2}^{\alpha,I} - \hat{q}_{p_2}^I, \dots, \hat{q}_{p_M}^{\alpha,I} - \hat{q}_{p_M}^I].$$

The left-hand equality of (15) can be written in compact form as  $\mathbf{Q}\hat{e} = e$ , where

$$\mathbf{Q} := \begin{bmatrix} \Phi_{p_1}(z_1) & \Phi_{p_2}(z_1) & \dots & \Phi_{p_{m_I}}(z_1) \\ \Phi_{p_1}(z_2) & \Phi_{p_2}(z_2) & \dots & \Phi_{p_{m_I}}(z_2) \\ \vdots & \vdots & \ddots & \vdots \\ \Phi_{p_1}(z_M) & \Phi_{p_2}(z_M) & \dots & \Phi_{p_{m_I}}(z_M) \end{bmatrix}. \tag{17}$$

The spectral coefficients  $\hat{e} = \mathbf{Q}^{-1}e$  are then jointly Gaussian,  $\hat{e} \sim \mathcal{N}(\mathbf{0}, \Sigma)$  with covariance matrix  $\Sigma := \sigma^2 \mathbf{Q}^{-1}(\mathbf{Q}^{-1})^\top = \sigma^2(\mathbf{Q}^\top \mathbf{Q})^{-1}$ , where each entry in the matrix  $\mathbf{Q}^\top \mathbf{Q}$  can be written as

$$[\mathbf{Q}^\top \mathbf{Q}]_{i,j} = \sum_{k=1}^M \Phi_{p_i}(z_k) \Phi_{p_j}(z_k).$$

Next, we reinterpret  $\frac{1}{M}[\mathbf{Q}^\top \mathbf{Q}]_{i,j}$  as a quadrature rule over  $\Gamma$  for approximating  $\left(\Phi_{p_i}, \Phi_{p_j}\right)_{L^2(\Gamma)}$  and exploit the orthonormality of  $\{\Phi_{p_i}\}_{i=1}^M$  to obtain

$$\frac{1}{M}[\mathbf{Q}^\top \mathbf{Q}]_{i,j} = \frac{1}{M} \sum_{k=1}^{m_I} \Phi_{p_i}(z_k) \Phi_{p_j}(z_k) \approx \left(\Phi_{p_i}, \Phi_{p_j}\right)_{L^2(\Gamma)} = \delta_{ij}, \tag{18}$$

where  $\delta_{ij}$  is the classical Kronecker delta. This entails that  $\mathbf{Q}^\top \mathbf{Q}$  is approximately diagonal,  $\mathbf{Q}^\top \mathbf{Q} \approx M \mathbb{1}$ , and thus  $\Sigma = \sigma^2(\mathbf{Q}^\top \mathbf{Q})^{-1} \approx \frac{1}{M} \sigma^2 \mathbb{1}$ . Therefore, the errors on the spectral coefficients are approximately statistically independent with variance  $\hat{\sigma}^2 = \frac{1}{M} \sigma^2$ . This implies that their absolute value approximately follows a folded normal distribution, and hence

$$\mathbb{E} \left[ \left| \hat{q}_p^{\alpha,I} - \hat{q}_p^I \right| \right] = \hat{\sigma} \sqrt{2/\pi} \tag{19}$$

for all  $p \in G(I)$ . This results in a spectral plateau in the decay of  $\hat{q}_p^{\alpha,I}$ .

We will later observe a spectral plateau to be present in more general scenarios than that of Example 1, cf. the numerical experiments of Sections 3 and 4. In Section 3 we present a method to detect the resulting spectral plateau effect, and embed this plateau detection into the adaptive MISC algorithm to formulate an automatic strategy named PLATEAUMISC, that is robust to intrinsic solver noise.

### 3. PlateauMISC: a multi-fidelity algorithm robust to solver errors

In this section we present the novel PLATEAUMISC algorithm. This builds upon the classic MISC algorithm which is detailed in Section 3.1. A plateau detection algorithm is described in Section 3.3 and the fully automatic PLATEAUMISC algorithm is presented in Section 3.4. An illustrative example highlights the advantages of the proposed PLATEAUMISC algorithm in Section 3.5.

### 3.1. Multi-index stochastic collocation

Multi-index stochastic collocation (MISC) is a multi-fidelity extension of the sparse grid stochastic collocation combination technique (4). The MISC algorithm has been investigated in [18,25,61–63] and applications have been investigated in [30,31,64]. Similar approaches (multi-level stochastic collocation) for parametric PDEs are considered in [21,65].

For a quantity of interest  $q : \Gamma \rightarrow \mathbb{R}$ , consider a hierarchy of computable fidelities  $q^\alpha(\mathbf{y}) \approx q(\mathbf{y})$  indexed by multi-indices  $\underline{\alpha} \in \mathbb{N}^n$ . Although a formal assumption is not made, it is expected that the fidelities become more accurate as the multi-index  $\underline{\alpha}$  increases in any dimension. For example, one could consider a sequence of solvers for an elliptic PDE problem given in domain  $\Omega \subset \mathbb{R}^n$  where mesh edge sizes are proportional to  $2^{-\alpha_i}$  for each spatial dimension  $i = 1, 2, \dots, n$ . Similarly, one could consider a time-dependent problem in which adaptive timestepping is applied with solvers using sequentially smaller local error tolerances  $\delta \propto 10^{-\alpha}$  as the fidelity index  $\alpha$  increases ( $n = 1$  in this case). We note that MISC does not directly apply when a hierarchical ordering of the fidelities cannot be established a priori (i.e., when increasing  $\alpha$  does not necessarily lead to a reduction of the error in the evaluation of a quantity of interest). In such cases, it is necessary to employ methods that invest some computational effort to identify correlations and effective hierarchies among the different fidelities; see, for example, [66,67].

Consider an admissible multi-index set  $I \subset \mathbb{N}^{n+d}$  defined over both the fidelity and parameter dimensions. A multi-index  $\underline{\nu} = [\underline{\alpha}, \underline{\beta}] \in I$  splits with  $\underline{\alpha} \in \mathbb{N}^n$  controlling the fidelity  $q^\alpha$  and  $\underline{\beta} \in \mathbb{N}^d$  again defining an interpolation operator to give a term  $I^{\underline{\beta}}[q^\alpha]$ . The combination technique approximation (4) is extended to give the MISC approximation,

$$q(\mathbf{y}) \approx q^{MISC}(\mathbf{y}) := I^I[q](\mathbf{y}) := \sum_{\underline{\nu} \in I} c_{\underline{\nu}} I^{\underline{\beta}}[q^\alpha](\mathbf{y}) \text{ for all } \mathbf{y} \in \Gamma \text{ where } \underline{\nu} = [\underline{\alpha}, \underline{\beta}]. \tag{20}$$

The resulting approximation is still a linear combination of global polynomial interpolants approximating the quantity of interest, with each interpolant acting on a different fidelity.

**Remark 1** (MISC is not interpolatory). Note that the resulting approximation (20) is not necessarily an interpolant: a collocation point  $\mathbf{z}$  may be requested by two interpolants acting on different fidelities  $\underline{\alpha}_1 \neq \underline{\alpha}_2$  with  $q^{\underline{\alpha}_1}(\mathbf{z}) \neq q^{\underline{\alpha}_2}(\mathbf{z})$ , yet the approximation  $q^{MISC}(\mathbf{z})$  can only take a single value.

The greedy adaptive construction of the MISC approximation (20) is presented in Algorithm 1 [18,30,61], based upon the classical Gerstner–Griebel-type dimensional-adaptive loop [28]. In particular, the construction follows [30, Algorithm 1]. Algorithm 1 is applied using a set of solvers giving approximations  $\{q^\alpha\}_{\alpha \in \mathbb{N}^n}$  where  $\alpha$  defines the fidelity. It iteratively grows the multi-index set  $I$  and consequently grows the polynomial approximation space  $\mathcal{P}^I(\Gamma; \mathbb{R})$ . Loosely speaking, adding a multi-index  $\underline{\nu} = [\underline{\alpha}, \underline{\beta}]$  in which we increase  $\underline{\alpha}$  refines the fidelity used by an interpolant defined via  $\underline{\beta}$ , whilst increasing  $\underline{\beta}$  will increase the number of evaluations and parametric approximation space for a particular fidelity  $\underline{\alpha}$ .

---

**Algorithm 1** Multi-index stochastic collocation.

---

```

1: function MISC( $\{q^\alpha : \Gamma \subset \mathbb{R}^d \rightarrow \mathbb{R}\}_{\alpha \in \mathbb{N}^n}$ )
2:   Initial multi-index set  $I = \{[\underline{1}, \underline{1}]\} \subset \mathbb{N}^{n+d}$ .
3:   while stopping criteria not met do
4:     Compute approximation  $q^{MISC} := \sum_{\underline{\nu} \in I} c_{\underline{\nu}} I^{\underline{\beta}}[q^\alpha]$  where  $\underline{\nu} = [\underline{\alpha}, \underline{\beta}]$ . ▷ Solve
5:     Compute reduced margin  $R_I$  via Eq. (8).
6:     for  $\underline{\nu} \in R_I$  do ▷ Estimate
7:       Estimate multi-index contribution  $E_{\underline{\nu}}$  via Eq. (21).
8:       Estimate multi-index cost  $W_{\underline{\nu}}$  via Eq. (22).
9:       Compute multi-index profit  $\mathfrak{p}_{\underline{\nu}} = E_{\underline{\nu}}/W_{\underline{\nu}}$ .
10:    end for
11:    Select  $\underline{\nu}^* \leftarrow \operatorname{argmax}_{\underline{\nu} \in R_I} \mathfrak{p}_{\underline{\nu}}$  and update  $I \leftarrow I \cup \underline{\nu}^*$ . ▷ Mark and Refine
12:  end while
13:  return  $q^{MISC} := \sum_{\underline{\nu} \in I \cup R_I} c_{\underline{\nu}} I^{\underline{\beta}}[q^\alpha]$  where  $\underline{\nu} = [\underline{\alpha}, \underline{\beta}]$ .
14: end function

```

---

At each iteration, the neighbouring multi-indices in the reduced margin  $R_I$  (8) are considered for refining the current multi-index set  $I$ . Error indicators  $E_{\underline{\nu}}$  are computed for each  $\underline{\nu} \in R_I$ . To define these, first it is assumed that for each parameter  $y_i$  there is a corresponding weight function  $\rho_i : \Gamma \rightarrow \mathbb{R}$  for  $i = 1, 2, \dots, d$ . For simplicity these are assumed to be  $\rho_i \equiv 1$  for  $i = 1, 2, \dots, d$ , i.e. each parameter can be considered to be representative of the image of an independent, uniformly distributed random variable on  $[0, 1]$ . Extension to other commonly used parameter domains and associated weight functions is straightforward and well established in the UQ literature [1–3]. The error indicators are defined as

$$E_{\underline{\nu}} := \left| \mathbb{E}[I^{I \cup \{\underline{\nu}\}}[q](\mathbf{y})] - \mathbb{E}[I^I[q](\mathbf{y})] \right|$$

$$:= \left| \int_{\Gamma} I^{I \cup \{\underline{\nu}\}}[q](\mathbf{y}) \rho(\mathbf{y}) d\mathbf{y} - \int_{\Gamma} I^I[q](\mathbf{y}) \rho(\mathbf{y}) d\mathbf{y} \right| \quad \text{where } \rho(\mathbf{y}) = \prod_{i=1}^d \rho_i(y_i). \tag{21}$$

Eq. (21) characterises the change in the approximation due to the addition of a multi-index  $\underline{v}$ . Alternative estimators can be used, for example considering the  $L^2_\rho(\Gamma; \mathbb{R})$  or  $L^\infty_\rho(\Gamma; \mathbb{R})$  norm of the difference  $I^{I \cup \{\underline{v}\}}[q] - I^I[q]$ . Next, for each fidelity  $\underline{\alpha}$ , a representative fidelity solve cost  $\widehat{W}_\alpha$  for a single approximation is introduced. The cost for adding a multi-index  $\underline{v} = [\underline{\alpha}, \beta]$  to the multi-index set  $I$  is the number of additional solves  $m_{I \cup \{\underline{v}\}}|_\alpha - m_I|_\alpha$  (i.e., the number of additional sparse grid points using fidelity  $\alpha$ , assuming nestedness of the one-dimensional collocation points  $\{Z^\beta\}_{\beta \in \mathbb{N}}$ ) multiplied by the representative fidelity solve cost  $\widehat{W}_\alpha$ , that is

$$W_{\underline{v}} := \widehat{W}_\alpha (m_{I \cup \{\underline{v}\}}|_\alpha - m_I|_\alpha). \tag{22}$$

The error indicator and cost are then combined to give a profit  $\mathfrak{p}_{\underline{v}}$  for each  $\underline{v} \in R_j$ , namely,

$$\mathfrak{p}_{\underline{v}} := E_{\underline{v}}/W_{\underline{v}}. \tag{23}$$

The greedy method in Algorithm 1 selects the largest profit and refines the approximation using the corresponding multi-index  $\underline{v}^*$ . The process repeats until a stopping criterion, for example a maximum total cost, maximum number of solves or a threshold profit, is reached. The final returned approximation uses the multi-index set  $I \cup R_j$ . This has no additional cost because the evaluations required to include  $R_j$  have already been performed when computing the error indicators.

### 3.2. Effect of solver error in MISC

The solver error discussed in Section 2.1 can lead the greedy approach in Algorithm 1 to mistakenly prioritise “noisy” fidelities during the refinement process. Each refinement step effectively adds higher-degree polynomial terms to the approximation space  $\mathcal{P}^I(\Gamma; \mathbb{R})$ . In the interpolation basis, these higher-degree terms may begin to fit the solver error rather than the true solution; equivalently, in the spectral polynomial basis, this refinement may introduce spectral terms with spurious, large coefficients, particularly in the spectral plateau, corresponding to the term (III) in the splitting (13). This leads to an inflated error indicator (21), causing the greedy Algorithm 1 to erroneously “chase the noise”: it will tend to overfit the solver error, being unable to identify the *true* solution behaviour. As a result, the adaptive procedure may overuse inaccurate, noisy (low fidelity) evaluations with falsely large error indicators, while overlooking more accurate fidelities. This behaviour can corrupt the approximation, resulting in a noisy and highly oscillatory surrogate  $q^{MISC}(y)$ . An example of this phenomenon will be discussed later in Section 3.5.

### 3.3. Plateau detection

To detect solver noise in the MISC approximation, we introduce a plateau detection algorithm. It exploits the truncated spectral polynomial expansion (12) and the expected spectral coefficient decay (14) to detect the effect of solver error via the spectral coefficient plateau resulting from term (III) in (13). The plateau detection borrows ideas from the approximation package CHEBFUN and its standardChop algorithm [[68], Table 3.3] [69]. Similar ideas have also been used in the context of shock capturing, where elementwise spectral-type polynomial expansions are used to detect non-smooth solutions [70,71].

We begin our discussion with the single-fidelity case before extending to the multi-fidelity setting. Consider the expansion (12) of the single fidelity approximation  $q^{\underline{\alpha}, I}(y)$  for a given fidelity level  $\underline{\alpha}$ . For an approximation using the multi-index set  $I$ , there is a corresponding set of polynomial degrees  $G(I)$  and corresponding computable spectral polynomial coefficients  $\{q_p^{\underline{\alpha}, I}\}_{p \in G(I)}$  defining the truncated spectral polynomial expansion. Let  $k_e := \max_{p \in G(I)} \|p\|_1$  be the maximum polynomial total degree in the truncated spectral polynomial expansion. We define a *spectral envelope*  $e : K \subset \mathbb{N}_0 \rightarrow \mathbb{R}$  with  $K := \{0, 1, \dots, k_e\}$  to provide an upper bound for the magnitude of the spectral coefficients, specifically

$$e(i) := \max_{p \in G(I) \text{ s.t. } \|p\|_1 \geq i} |q_p^{\underline{\alpha}, I}|. \tag{24}$$

The envelope  $e$  is expected to show two regimes: an exponential decrease as predicted by the a priori bound (14), followed by a spectral plateau due to solver error as predicted by (III) in (13). This behaviour will be illustrated in the forthcoming numerical examples. To identify this plateau, we fit a piecewise linear model to  $\log_{10}(e)$  and detect flattening in the tail. Note that the a priori decay rate (14) may not accurately describe the low polynomial degree terms. Therefore, we ignore the initial  $k_{in}$  entries of the envelope (*burn in*). Similarly, the highest polynomial degree spectral coefficients may be poorly estimated due to aliasing effects (i.e., regime (II) in splitting (13)) so we discard the final  $k_{out}$  terms (*burn out*) in the plateau detection process.

In practice, a *change point detection* algorithm with a piecewise linear model is used to partition  $\{(i, \log_{10}(e(i)))\}_{i=k_{in}}^{k_e-k_{out}}$  into two with the change point  $j$  [72]. These fitted functions will later be referred to as *piecewise log-linear* functions. The models are defined by minimising the sum of the square errors,

$$\begin{aligned} \{j, k_0, c_0, k_1, c_1\} = & \operatorname{argmin}_{j \in \{k_{in}, \dots, k_e - k_{out}\}} \sum_{i=k_{in}}^{j-1} (\log_{10}(e(i)) - m_0 i - c_0)^2 \\ & + \sum_{i=j}^{k_e - k_{out}} (\log_{10}(e(i)) - m_1 i - c_1)^2. \end{aligned} \tag{25}$$

The change point  $j$  is returned along with two linear models

$$\begin{cases} m_0 i + c_0, & i \in \{k_{in}, \dots, j-1\}, \\ m_1 i + c_1, & i \in \{j, \dots, k_e - k_{out}\}. \end{cases} \quad (26)$$

The second model is associated with the spectral plateau. As already hinted, our algorithm identifies a plateau if the following conditions are met:

1. the second linear model has a small gradient  $|m_1| \leq m_*$ ,
2. the second model is suitably long  $k_e - k_{out} - j > k_{min}$ .

The spectral plateau level is computed as  $\mathcal{E}^{plateau} := 10^{m_1 j + c_1}$ . Spectral plateau detection is summarised in Algorithm 2. In practice, we use the optimised MATLAB `findchangepts` solver to identify the models in (25) and (26). Although ideas are drawn from the standardChop algorithm in [69], the standardChop algorithm is not appropriate in our context because it is designed for identification of spectral plateau effects at or near numerical precision.

Four hyperparameters have been introduced in the plateau detection algorithm detailed above ( $k_{in}$ ,  $k_{out}$ ,  $m_*$ ,  $k_{min}$ ). The following considerations can be useful when tuning them:

- Burn in  $k_{in}$ : This should be tuned based upon the expected decay of the spectral coefficients of the target response surface,  $\hat{q}_p$ , in the ideal case in which the solver noise is absent (see Eq. (13)). If the surface is expected to be extremely smooth it can be set equal to zero. If conversely the surface is expected to have e.g. local features, steep gradients, etc. there will be a pre-asymptotic behaviour that must be captured by the spectral polynomial expansion before the exponential decay in coefficient magnitude is observed. From our experiments, we suggest a burn in of  $k_{in} = 2$ .
- Burn out  $k_{out}$ : This is tuned to remove any aliasing effects in the tail of the spectral envelope. An example of such effects will be presented in the numerical results of Section 4.1. From our numerical experiments, we suggest a burn out of  $k_{out} = 2$ .
- Gradient threshold  $m_*$ : A gradient threshold of zero is almost impossible to attain in practice, whilst a gradient of  $O(1)$  would still indicate spectral coefficient decay. We suggest a gradient threshold of  $m_* = 0.1$ , which could be decreased if the response surface is expected to experience a fast coefficient decay.
- Minimum plateau length  $k_{min}$ : The minimum plateau length must be chosen to balance the benefit of ensuring all useful information is extracted from cheaper, low fidelity simulations, whilst avoiding unnecessary computational burden oversampling the spectral plateau. We suggest a minimum plateau length  $k_{min} = 3$ . This could be increased if higher fidelity solvers are significantly more expensive, making it crucial to extract as much information as possible from the cheaper solvers.

It is worth noting that the hyperparameter values are problem-dependent. Nonetheless, the PlateauMISC strategy exhibited behaviour that was robust to their selection. The numerical experiments in Section 3 and Section 4 followed these principles to determine suitable values.

---

**Algorithm 2** Spectral plateau detection.

---

```

1: function PLATEAUDETECTION( $\{i, e(i)\}_{i=0}^{k_e}, k_{in}, k_{out}, k_{min}$ )
2:   Compute  $[j, m_0, c_0, m_1, c_1]$  via (25)
3:   if  $|m_1| < m_*$  and  $k_e - k_{out} - j > k_{min}$  then
4:      $\mathcal{E}^{plateau} \leftarrow 10^{m_1 j + c_1}$ 
5:      $isPlateau \leftarrow \text{true}$ 
6:   else
7:      $\mathcal{E}^{plateau} \leftarrow 0$ 
8:      $isPlateau \leftarrow \text{false}$ 
9:   end if
10:  return  $j, \mathcal{E}^{plateau}, isPlateau$ 
11: end function

```

---

### 3.4. PLATEAUMISC variant of MISC

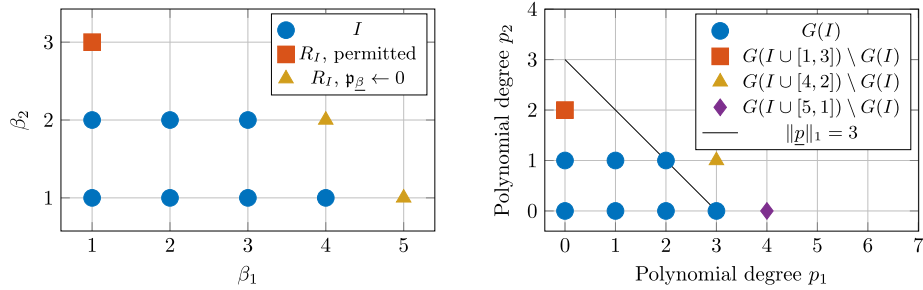
The plateau detection procedure described in Algorithm 2 is now integrated into the adaptive MISC approximation framework of Algorithm 1. To achieve this, Algorithm 1 is modified by incorporating a detect step in the standard solve-estimate-mark-refine loop, enabling the identification of the spectral plateaus. Since MISC is a multi-fidelity approximation method, plateau detection is applied independently at each fidelity level.

Consider an approximation  $q^{MISC}$  of the form (20). Denote the set of active fidelities in the approximation by  $A$ , that is,

$$A = \{\underline{\alpha} \text{ such that there exists a } \underline{\beta} \text{ for which } [\underline{\alpha}, \underline{\beta}] \in I\}. \quad (27)$$

For each active fidelity  $\underline{\alpha} \in A$ , a restriction to the fidelity is formed:

$$q^{MISC}|_{\underline{\alpha}} := I|_{\underline{\alpha}}[q] \text{ where } I|_{\underline{\alpha}} := \left\{ \hat{\beta} \text{ s.t. } [\hat{\underline{\alpha}}, \hat{\beta}] \in I \text{ s.t. } \hat{\underline{\alpha}} = \underline{\alpha} \right\}. \quad (28)$$



(a) Multi-index set  $I$ , reduced margin  $R_I$  and subset of reduced margin with profits  $p_{\beta}$  set to zero. (b) Polynomial space  $G(I)$  corresponding to the multi-index set  $I$ , and the additional polynomials  $G(I \cup \underline{p}) \setminus G(I)$  added by the multi-indices  $\underline{p} \in R_I$ .

Fig. 4. Example demonstrating the filtering of profits specified in (29). An example multi-index set  $I$  is shown in Fig. 4a. The polynomial space  $G(I)$  is shown in Fig. 4b when using  $I$  with the ‘linear’ level-to-knots rule from (2). If the corresponding fidelity is saturated with change point  $j = 3$ , the profits for multi-indices  $[4, 2]$  and  $[5, 1]$  are set to zero because they only add polynomial terms of total degree  $\|p\|_1 \geq 3$ . We still permit the multi-index  $[1, 3]$  because it adds at least one polynomial term of total degree  $\|p\|_1 < 3$ .




Fig. 5. Modified reduced margin  $R_{I,S}$  defined in (31) for a multi-index set with  $n = 1$  and  $d = 1$ . If the saturated fidelities are  $S = \{1, 2\}$ , the multi-index  $[3, 3]$  is still permitted to be explored by the adaptive algorithm, whilst the multi-indices  $[1, 3]$ ,  $[2, 3]$  may have their profits set to zero by (29). To maintain admissibility the multi-indices  $B_{I,[3,3]}$ , as defined in (32), must be used to ‘backfill’ the multi-index set if  $[3, 3]$  is selected by the adaptive algorithm.

Each multi-index set  $I|_{\alpha}$  is downward-closed by construction, therefore (28) defines valid single-fidelity SGSC approximations. For each restricted approximation  $q^{MISC}|_{\alpha}$ , an envelope  $e|_{\alpha}$  can be defined via (24), and change points  $j|_{\alpha}$  and plateau levels  $\mathcal{E}^{plateau}|_{\alpha}$  can be computed using Algorithm 2. Fidelities that have been identified as having a spectral plateau are collected as the set of saturated fidelities  $S \subset \mathbb{N}^n$ . High polynomial degree terms for saturated fidelities should not be marked for further parametric refinement, limiting their employment in the construction of the surrogate model.

To stop the greedy algorithm from adding high polynomial degree terms for the saturated fidelities  $S$ , the mark-refine step in Algorithm 1 is preceded by setting a subset of profits to zero. This subset corresponds to multi-indices  $\underline{v} = [\underline{\alpha}, \underline{\beta}] \in R_I$  that satisfy the following criteria (see also Fig. 4):

- $\underline{v}$  corresponds to a saturated fidelity  $\alpha \in S$ ;
- $\underline{v}$  only adds polynomial terms  $\underline{p} \in G(I|_{\alpha} \cup \underline{\beta}) \setminus G(I|_{\alpha})$  of polynomial total degree greater than or equal to the fidelity’s change point  $j|_{\alpha}$ , that is if  $\|p\|_1 \geq j|_{\alpha}$  for all  $\underline{p} \in G(I|_{\alpha} \cup \underline{\beta}) \setminus G(I|_{\alpha})$ .

The second point allows low polynomial total degree terms to still be added via refinement, even for saturated fidelities. These conditions can be expressed as

$$p_{\underline{v}} \leftarrow 0 \text{ if } \underline{\alpha} \in S \text{ and } \|p\|_1 \geq j|_{\alpha} \text{ for all } \underline{p} \in G(I|_{\alpha} \cup \underline{\beta}) \setminus G(I|_{\alpha}) \quad (29)$$

where  $\underline{v} = [\underline{\alpha}, \underline{\beta}]$  and  $j|_{\alpha}$  is the change point for the associated fidelity  $\alpha$  computed by Algorithm 2.

Another technical challenge arises from the treatment of the saturated fidelities  $S$ . Blocking further refinement of these saturated fidelities may unintentionally restrict access to high-degree polynomial approximations at unsaturated fidelity levels. This issue stems from the admissibility constraint on the multi-index set, which is enforced by only exploring the reduced margin  $R_I$  as defined in equation (8). For instance, consider Fig. 5: the multi-index  $\underline{v} = [3, 3]$  may not be reached if the fidelities  $\alpha = 1, 2$  are already saturated. However, fidelity  $\alpha = 3$  being less noisy, may still benefit from further parametric refinement.

To enable this, we introduce a modified reduced margin that permits selective refinement at unsaturated fidelities: first define the margin set  $M_I$ ,

$$M_I := \left\{ \underline{\mu} \in \mathbb{N}^{n+d} \text{ s.t. } \underline{\mu} = \underline{v} + \underline{e}_i \text{ for some } \{\underline{v}, i\} \in I \times \{1, 2, \dots, n + d\} \right\} \quad (30)$$

then the *modified reduced margin* is

$$R_{I,S} := \left\{ \underline{v} \in M_I \text{ s.t. } [\underline{\alpha}, \underline{\beta}] := \underline{v} - \underline{e}_i \in I \text{ for all } i = 1, 2, \dots, n+d \text{ or } \underline{\alpha} \in S \right\}. \quad (31)$$

This construction ensures that each backward neighbour  $\underline{v} - \underline{e}_i$  either belongs to the current multi-index set  $I$  or corresponds to a saturated fidelity in  $S$  (i.e.  $\underline{v} - \underline{e}_i \in \bigcup_{\underline{\alpha} \in S} \{[\underline{\alpha}, \underline{\beta}]\}$  for all  $\underline{\beta} \in \mathbb{N}^d$ ). Multi-indices corresponding to high polynomial degree terms in saturated fidelities are not explicitly removed from the reduced margin, however they are ignored in the adaptive algorithm since their profits are set to zero by (29).

To maintain admissibility of the multi-index set when adding a multi-index from the modified reduced margin  $R_{I,S}$ , we introduce a backfilling procedure. For  $\underline{v} \in R_{I,S}$  the indices required to backfill the multi-index set  $I$  to maintain admissibility are denoted by

$$B_{I,\underline{v}} := \{ \text{minimal set of } \hat{\underline{v}} \in \mathbb{N}^{n+d} \text{ s.t. } I \cup B_{I,\underline{v}} \cup \{ \underline{v} \} \text{ satisfies admissibility property(7)} \}, \quad (32)$$

see again Fig. 5 for a visualisation. By construction all multi-indices  $\underline{v} \in B_{I,\underline{v}}$  can only correspond to saturated fidelities. The refinement of the multi-index set  $I$  with  $\underline{v} \in R_{I,S}$  is then updated as  $I \leftarrow I \cup B_{I,\underline{v}} \cup \{ \underline{v} \}$ . Finally, the introduction of backfilling requires that we update the definitions of the error indicator  $E_{\underline{v}}$  and cost  $W_{\underline{v}}$  for a multi-index  $\underline{v} \in R_{I,S}$  accounting for the backfill set  $B_{I,\underline{v}}$ . The error indicator (21) is thus updated to

$$E_{\underline{v}} := \left| \mathbb{E} \left[ \mathcal{I}^{I \cup B_{I,\underline{v}} \cup \{ \underline{v} \}} [q] \right] - \mathbb{E} \left[ \mathcal{I}^I [q] \right] \right|, \quad (33)$$

and similarly the cost (22) is updated to

$$W_{\underline{v}} := \widehat{W}_{\underline{\alpha}}(m_{I_{\underline{\alpha}} \cup \{ \underline{\beta} \}} - m_{I_{\underline{\alpha}}}) + \sum_{\hat{\underline{v}} \in \{ \hat{\underline{\alpha}}, \hat{\underline{\beta}} \} \cap B_{I,\underline{v}}} \widehat{W}_{\hat{\underline{\alpha}}}(m_{I_{\hat{\underline{\alpha}} \cup \{ \hat{\underline{\beta}} \}} - m_{I_{\hat{\underline{\alpha}}}}) \quad (34)$$

where the cost required to backfill the multi-index set must be included.

---

### Algorithm 3 Plateau detection multi-index stochastic collocation.

---

```

1: function PLATEAUMISC( $\{q^{\underline{\alpha}} : \Gamma \subset \mathbb{R}^d \rightarrow \mathbb{R}\}_{\underline{\alpha} \in \mathbb{N}^n}$ )
2:   Initial multi-index set  $I = \{[\underline{1}, \underline{1}]\} \subset \mathbb{N}^{n+d}$ 
3:   while stopping criteria not met do
4:     Compute approximation  $q^{MISC} := \sum_{\underline{v} \in I} c_{\underline{v}} \mathcal{I}^{\underline{v}} [q^{\underline{\alpha}}]$  where  $\underline{v} = [\underline{\alpha}, \underline{\beta}]$ . ▷ Solve
5:     Compute modified reduced margin  $R_{I,S}$  via Eq. (31).
6:     for  $\underline{v} \in R_{I,S}$  do ▷ Estimate
7:       Estimate multi-index contribution  $E_{\underline{v}}$  via Eq. (33).
8:       Estimate multi-index cost  $W_{\underline{v}}$  via Eq. (34).
9:       Compute multi-index profits  $\mathfrak{p}_{\underline{v}} = E_{\underline{v}} / W_{\underline{v}}$ .
10:    end for
11:
12:    for each unsaturated active fidelity  $\underline{\alpha} \in A \setminus S$  do ▷ Detect
13:      Construct  $q^{MISC}|_{\underline{\alpha}}$  via (28).
14:      Construct spectral envelope  $(i, e|_{\underline{\alpha}}(i))$  via (24).
15:       $j|_{\underline{\alpha}}, \mathcal{E}^{plateau}|_{\underline{\alpha}}, isPlateau|_{\underline{\alpha}} = \text{PLATEAUDETECTION}(i, e|_{\underline{\alpha}}(i), k_{in}, k_{out}, k_{min})$ 
16:    end for
17:    Update saturated fidelities  $S \leftarrow S \cup \{ \underline{\alpha} \in A \setminus S \text{ s.t. } isPlateau|_{\underline{\alpha}} = \text{true} \}$ 
18:
19:    Set zero profits  $\mathfrak{p}_{\underline{v}}$  for saturated  $\underline{v}$  using criteria in (29).
20:    Select  $\underline{v}^* \leftarrow \text{argmax}_{\underline{v} \in R_{I,S}} \mathfrak{p}_{\underline{v}}$  and refine  $I \leftarrow I \cup B_{I,\underline{v}^*} \cup \{ \underline{v}^* \}$ . ▷ Mark and Refine
21:  end while
22:  return  $q^{MISC} := \sum_{\underline{v} \in I} c_{\underline{v}} \mathcal{I}^{\underline{v}} [q^{\underline{\alpha}}]$  where  $\underline{v} = [\underline{\alpha}, \underline{\beta}]$ .
23: end function

```

---

The complete PLATEAUMISC algorithm is presented in Algorithm 3. Note the differences with respect to Algorithm 1:

- the plateau detection step is introduced to identify saturated fidelities;
- a subset of profits may be set to zero via (29);
- the returned approximation no longer uses the reduced margin.

This third point is to ensure that the high degree polynomial terms used to identify the spectral plateaus in the saturated fidelities are not included in the returned approximation, since they will generally only contribute spurious oscillations to the full approximation.

### 3.5. Illustrative validation example

The PLATEAUMISC algorithm is validated using an analytic two-dimensional Gaussian peak (2DGP) test function from the Genz suite of test functions [73]. Consider the function  $q : \Gamma := [0, 1]^2 \rightarrow \mathbb{R}$

$$q(\mathbf{y}) := \exp(-C_1^2(y_1 - 0.5)^2) \exp(-C_2^2(y_2 - 0.5)^2) \quad (35)$$

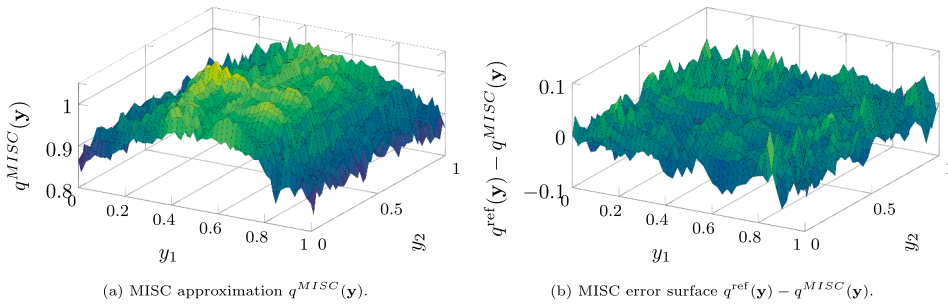


Fig. 6. MISC approximation  $q^{MISC}$  of the 2DGP Genz test function (35) using Algorithm 1 with fidelities  $q^\alpha$  defined in (36).

where  $C_1 = 2^{-2}C$ ,  $C_2 = 3^{-2}C$  and  $C = \frac{36}{13}$ . A controlled fidelity error is artificially introduced by considering a sequence of approximations of the quantity of interest  $q^\alpha : \Gamma \rightarrow \mathbb{R}$  for  $\alpha \in \mathbb{N}$ , such that

$$q^\alpha(\mathbf{y}) := q(\mathbf{y}) + 10^{-2\alpha}x(\mathbf{y}) \tag{36}$$

where, for each parameter  $\mathbf{y}$ , the solver error  $x(\mathbf{y})$  is a realisation of a standard normal random variable  $X \sim \mathcal{N}(0, 1)$ . Evaluations for each fidelity  $\alpha$  have a corresponding solve cost  $\widehat{W}_\alpha = 10^\alpha$ . This setup mimics a scenario in which the solver error decreases exponentially with fidelity level  $\alpha$ , while the fidelity evaluation cost increases exponentially. This cost–accuracy trade off is consistent with many numerical solvers; for example, in  $d$ -dimensional finite element methods, uniformly refining the mesh by a factor of  $10^\alpha$  increases the number of degrees of freedom, and hence the computational cost, by a factor of  $10^{ad}$ , while the approximation error decreases as  $10^{-\alpha p}$  for some  $p > 0$  given by the accuracy of the employed discretisation.

MISC approximations are constructed using symmetric Leja points [29] with the ‘two-step’ growth rule defined in (2). Errors are computed numerically with respect to a high fidelity reference solution. This reference  $q^{\text{ref}}$  is obtained via a single-fidelity sparse grid approximation using the combination technique (4) with Clenshaw–Curtis points and the doubling rule of (2), using the highest fidelity  $\alpha^{\text{ref}} = 8$  and an isotropic Smolyak multi-index set

$$I_{\text{ref}} = \{\underline{\beta} \text{ s.t. } \|\underline{\beta}\|_1 \leq 2 + w\} \text{ for } w = 8. \tag{37}$$

Intermediate approximations are computed at  $w = 0, 1, 2, \dots, 7$ , in addition to a single reference fidelity adaptive SC approximation. The use of reference fidelity  $\alpha^{\text{ref}} = 8$  ensures a reference solver error of the order of  $10^{-16}$  (i.e. near machine precision) and provides a reliable benchmark for evaluating parametric convergence and model performance.

The classical MISC response surface  $q^{MISC}$  obtained using Algorithm 1 has a total computational cost of 2891 and it is shown in Fig. 6a. This is clearly corrupted by the low fidelity solve error. The corresponding error surface is shown in Fig. 6b.

The  $L^2_\rho(\Gamma; \mathbb{R})$  approximation error is estimated as

$$\|q - q^{MISC}\|_{L^2_\rho(\Gamma; \mathbb{R})} \approx \frac{1}{|\Gamma|} \int_\Gamma |q^{\text{ref}}(\mathbf{y}) - q^{MISC}(\mathbf{y})|^2 d\mathbf{y} \approx \left( \frac{1}{N_{MC}} \sum_{i=1}^{N_{MC}} |q^{\text{ref}}(\mathbf{y}_i) - q^{MISC}(\mathbf{y}_i)|^2 \right)^{1/2} \tag{38}$$

where  $\rho(\mathbf{y}) = 1/|\Gamma| = 1$  is used as the weight function in the integral. Eq. (38) is computed using  $N_{MC} = 10^4$  Monte Carlo samples. The  $H^1_\rho(\Gamma; \mathbb{R})$  norm of the approximation error is estimated similarly, using a finite difference approximation for the gradient of  $q^{\text{ref}}$  and  $q^{MISC}$  as implemented in the SPARSE GRIDS MATLAB KIT [29]. The error plots in Fig. 7a and b show that the classical MISC approximation constructed by Algorithm 1 has not delivered an accurate approximation to the 2DGP function defined in (35) and the approximation quality has stagnated. Here, and throughout the experiments, computational cost will refer to the sum of evaluation cost for all of the evaluations required by the approximation algorithm. The approximation is stuck adding cheap, low fidelity evaluations. The algorithm eventually terminates, using 505 points at fidelity  $\alpha = 1$ , resulting in the surrogate shown in Fig. 6a.

The sampling profile shown in Fig. 8a shows the number of solves  $m|_\alpha$  for each fidelity against the cumulative computational work of the MISC algorithm. Each marker represents an iteration of the algorithm.

Many more grid points  $m|_\alpha$  are sampled for fidelity  $\alpha = 1$ , whilst only one grid point is sampled for fidelity  $\alpha = 3$ . By definition in (36), it is clear that the fidelity  $\alpha = 1$  is extremely erroneous. Fig. 8b shows the spectral polynomial coefficients  $\{\widehat{q}_p^{MISC}|_\alpha\}_{p \in G(I)}$  for each fidelity in the computed MISC approximation  $q^{MISC}$ . Notice both the fidelities  $\alpha = 1$  and  $\alpha = 2$  have notable spectral plateaus in the polynomial coefficients. For fidelity  $\alpha = 1$ , it is clear that the solver error is spoiling the computed spectral polynomial coefficients, and consequently also leading to an approximation of poor quality. The plateau is lowered by using a more accurate solver, with  $\alpha = 2$ . There is only one solve on the fidelity  $\alpha = 3$  so there is only the constant polynomial  $p = 0$  plotted in Fig. 8b.

### 3.5.1. Plateau detection

Plateau detection via Algorithm 2 is applied to the fidelities in the classical MISC approximation with  $k_{in} = k_{out} = 2$ . and  $k_{min} = 3$ . The envelopes for the three active fidelities  $\alpha = 1, 2, 3$  are shown in Fig. 9a. For  $\alpha = 1, 2$ , these capture the exponential decrease in coefficient magnitude and the plateau in the coefficient magnitude due to solver noise. The piecewise log-linear models are shown in Fig. 9b. The effect of the burn in  $k_{in}$  and burn out  $k_{out}$  are observed. The first models do not start at zero, but at polynomial total

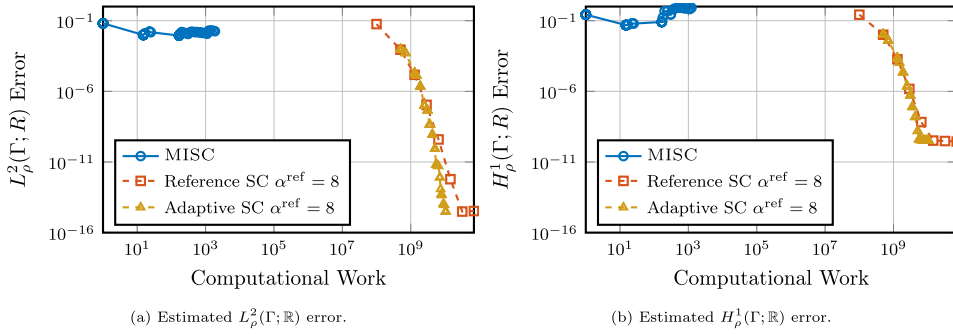


Fig. 7. Estimated  $L^2_p(\Gamma; \mathbb{R})$  and  $H^1_p(\Gamma; \mathbb{R})$  error for the MISC approximation  $q^{MISC}$  of the 2DGP Genz test function (36) using Algorithm 1 with fidelities  $q^\alpha$  defined in (36).

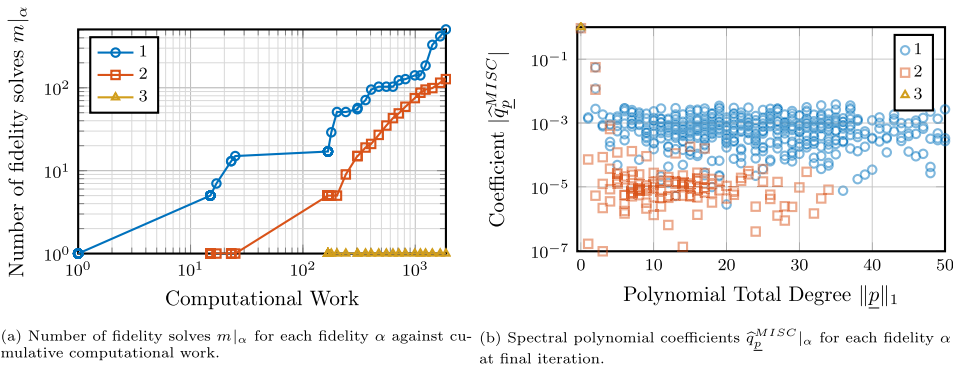


Fig. 8. Number of fidelity solves against computational work and snapshot of spectral polynomial coefficients for the MISC approximation  $q^{MISC}$  of the 2DGP Genz test function (36) using Algorithm 1 with fidelities  $q^\alpha$  defined in (36).

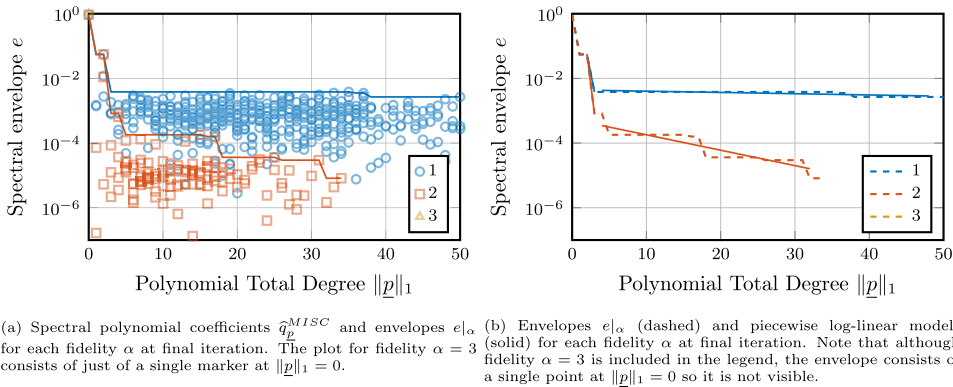


Fig. 9. Spectral polynomial coefficients, spectral envelopes and piecewise log-linear models for the MISC approximation  $q^{MISC}$  of the 2DGP Genz test function (36) using Algorithm 1 with fidelities  $q^\alpha$  defined in (36).

degree equal to two. Similarly, the second models finish two units before the end of the envelopes. The slopes of the second piecewise log-linear models are  $-8.0 \times 10^{-4}$  and  $-1.4 \times 10^{-2}$  for fidelities  $\alpha = 1, 2$  respectively. In both cases a plateau is successfully identified by Algorithm 2. The computed plateau levels are  $\mathcal{E}^{plateau} = 4.5 \times 10^{-3}$  and  $8.4 \times 10^{-5}$  respectively. This matches intuition: the solver error defined in (36) for fidelity  $\alpha = 1$  has standard deviation  $10^{-2}$  and for fidelity  $\alpha = 2$  has standard deviation  $10^{-4}$ . For fidelity  $\alpha = 3$  the envelope in Fig. 9b is not visible — it is a single point, corresponding to the single polynomial coefficient for  $\alpha = 3$ !

### 3.5.2. PLATEAUMISC

The PLATEAUMISC strategy in Algorithm 3, is now applied: comparing the resulting response surface (Fig. 10a) to the MISC surrogate in Fig. 6a it is clear that Algorithm 3 offers an improved response surface — it is far less noisy. The corresponding error surface  $q^{ref}(y) - q^{MISC}(y)$  provided by PLATEAUMISC is reported in Fig. 10b.

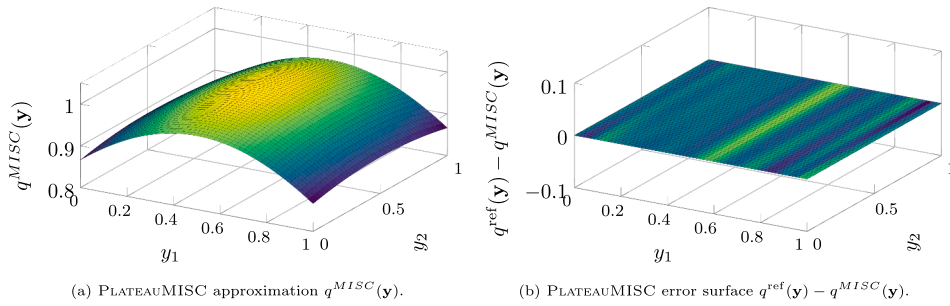


Fig. 10. PLATEAUMISC approximation  $q^{MISC}$  of the 2DGP Genz test function (35) using Algorithm 3 with fidelities  $q^\alpha$  defined in (36).

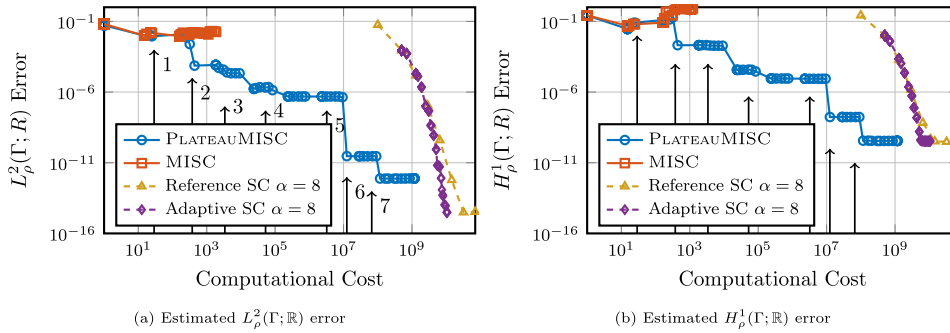


Fig. 11. Estimated  $L^2_\rho(\Gamma; \mathbb{R})$  and  $H^1_\rho(\Gamma; \mathbb{R})$  approximation error for the PLATEAUMISC approximation  $q^{MISC}$  of the 2DGP Genz test function (36) using Algorithm 3 with fidelities  $q^\alpha$  defined in (36). Vertical arrows denote the times at which a fidelity is marked as saturated.

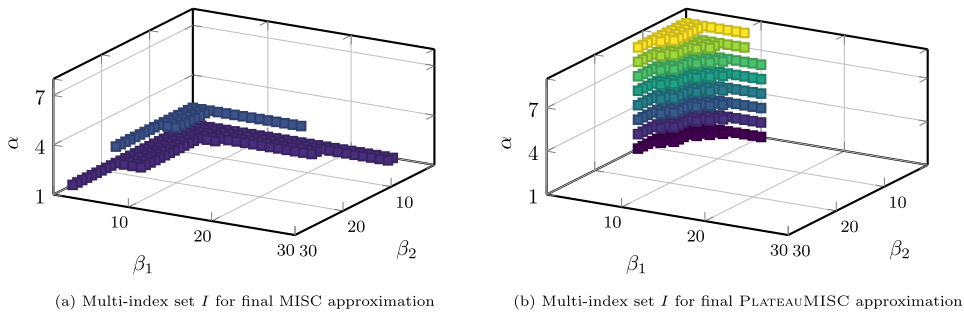
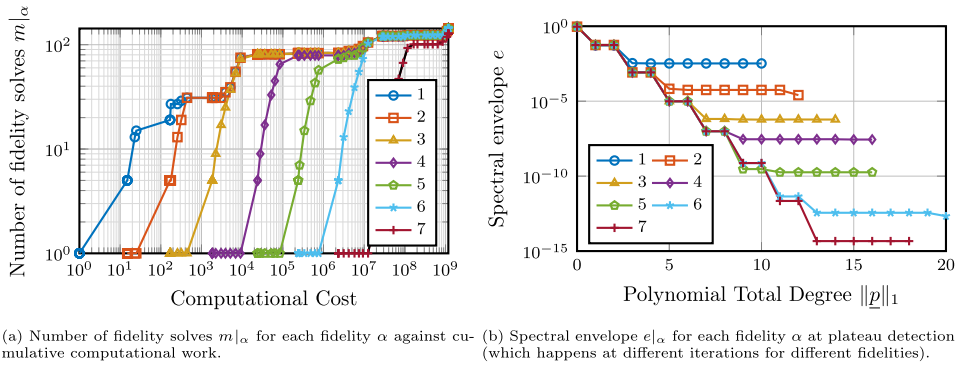


Fig. 12. Multi-index sets  $I$  for the final MISC and PLATEAUMISC approximations of the 2DGP Genz test function (36) using Algorithms 1 and 3 with fidelities  $q^\alpha$  defined in (36).

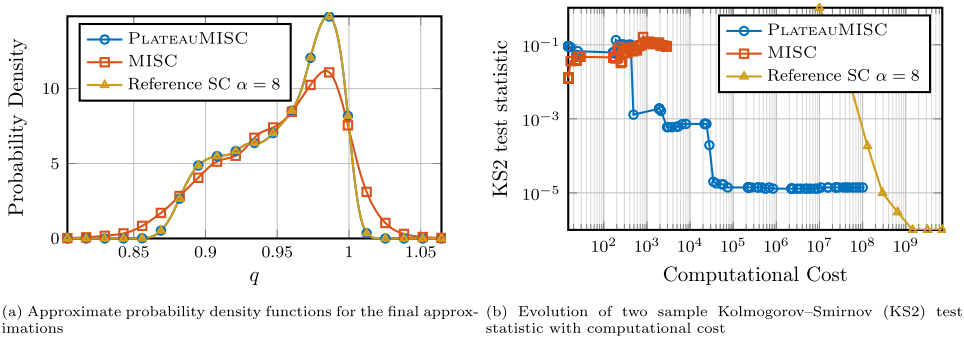
To make a fair comparison, the approximation error is plotted against computational cost for both approximations. The estimated  $L^2_\rho(\Gamma; \mathbb{R})$  and  $H^1_\rho(\Gamma; \mathbb{R})$  errors are shown in Fig. 11a and b.

The PLATEAUMISC approximation error decreases as we increase the computational budget since it can successfully transition to higher available fidelities, whereas the classical MISC approximation stagnates at a large error and cannot proceed further as it requests too many points from the noisy low fidelities. Vertical arrows in Fig. 11a (and later figures) indicate points at which a fidelity is marked as saturated in the PLATEAUMISC approximation. These are generally aligned with a decrease in the approximation error, however the approximation is not returned at every iteration of the algorithm (output iterations are controlled by computational budget) so there is a discrepancy at some points. The benefit of the multi-fidelity approach is also shown. For comparable computational cost, a more accurate approximation is constructed by PLATEAUMISC than the single high fidelity reference or a single reference fidelity adaptive stochastic collocation approximation.

The final multi-index sets are shown in Fig. 12: it is clear that MISC oversamples the low fidelities and does not access the more accurate fidelities with higher  $\alpha$ . The effect of backfilling in PLATEAUMISC is also clearly observed: the lower fidelity levels (smaller  $\alpha$ ) in the PLATEAUMISC multi-index set are enriched to allow the higher fidelities to access higher polynomial degree parametric approximation spaces. For problems with stronger correlations in the errors between fidelities, the multi-index set will be more anisotropic with respect to fidelity since more information can be captured by the low fidelities. This will be observed later in Section 4.1.



**Fig. 13.** Number of fidelity solves against computational work and snapshot of spectral polynomial coefficients for the PLATEAUMISC approximation  $q^{MISC}$  of the 2DGP Genz test function (36) using Algorithm 3 with fidelities  $q^\alpha$  defined in (36).



**Fig. 14.** Final approximate probability density functions and evolution of two sample Kolmogorov–Smirnov test statistic for the 2DGP Genz test function (36) using Algorithm 3 with fidelities  $q^\alpha$  defined in (36).

From the fidelity work allocations reported in Fig. 13a, we see that computational cost is saved in PLATEAUMISC as the high fidelities are sequentially activated, and only when the parametric approximation is sufficiently accurate for it to be worthwhile. The effect of backfilling is seen as the number of solves will continue to increase for all fidelities, despite some being marked as saturated and blocked from explicit refinement. Both the exponential decrease in the spectral coefficient magnitudes and the plateau effect from the solver error are clearly observed in the spectral polynomial coefficients in Fig. 13b.

Finally, we consider approximating the probability density function (PDF)  $\rho_q : \Gamma \rightarrow \mathbb{R}$  of the quantity of interest for uniformly distributed input parameters, that is estimating  $\rho_q$  such that

$$\text{Prob}(q \in A) = \int_A \rho_q(\mathbf{y}) d\mathbf{y}. \tag{39}$$

The PDF is estimated using  $N_{MC} = 10^4$  Monte Carlo samples from the surrogate models and the MATLAB `ksdensity` function.

Fig. 14a shows that the PLATEAUMISC PDF approximation closely matches the reference PDF approximation whereas the classical MISC PDF approximation is significantly less accurate. This is also supported by the two-sample Kolmogorov–Smirnov test statistic, as shown in Fig. 14b. This is produced using  $N_{MC} = 10^6$  samples from the surrogates and the reference approximation. Even at relatively low computation cost, the PLATEAUMISC approximation gives a good approximation of the PDF, whilst the classical MISC approximation shows no convergence with computational cost at all. The plateau in the KS2 test statistic is hypothesised to be a consequence of discretisation due to the finite sample sets.

#### 4. Numerical examples

To illustrate the practical relevance of the proposed PLATEAUMISC method we consider two representative test cases: an unsteady advection–diffusion problem with uncertain coefficients, and a parametric turbulent incompressible Navier–Stokes flow. The first example is a parametric variant of the well-known *double glazing problem*, originally introduced in [74,75] and based on the benchmark problems from the International Association for Hydraulic Research [76,77]. This problem features a two-dimensional domain in space, a time dependence, and two uncertain parameters, leading to a high-dimensional problem of dimension 5. The second example involves a parametric version of the *Two Dimensional Wall Mounted Hump (2DWMH)* problem, proposed by NASA for turbulence model validation [78,79]. This highly non-linear problem has been extended in [80] by introducing a parametric suction jet to control the




Fig. 15. Realisations of the advection field  $w(x, y)$  defined in (41) for three parameter realisations  $y = [-1, -1], [0, 0]$  and  $[1, 1]$ .

reattachment point of the separated flow downstream of the ‘hump’. In our work, the Reynolds number is also introduced as a second parameter, leading to a problem set in a high-dimensional space of dimension 4.

In both cases, the fidelity is controlled by a single index ( $n = 1$ ) associated with the accuracy of the solver employed for the computation of each instance of the parametric problem, while the parametric dimension is set to  $d = 2$ . It is worth noticing that MISC has already been successfully validated for problems with parametric dimension larger than  $d = 2$ , see e.g. [18,19,25], where experimental tests for up to tens, hundreds, and thousands of parameters were respectively presented. The scope of this work is instead to investigate the impact of the solver error on the accuracy of the surrogate model of the QoI and the development of an automatic strategy to make MISC robust in this framework. The presented numerical experiments thus focus on exploring different sources of such error and their interaction with nonlinearity, rather than on high-dimensional problems. We point out that the plateau detection algorithm is essentially dimension-independent in terms of cost, so its execution as parameter dimension  $d$  increases does not pose any additional challenge compared to the traditional versions of MISC.

#### 4.1. Uncertainty quantification for unsteady advection–diffusion problem with uncertain coefficients

Consider the spatial domain  $D := (-1, 1)^2$ , finite time  $T > 0$  and parameter domain  $\Gamma := [0, 1]^2$ . For each  $y \in \Gamma$ , we seek the solution to the following advection–diffusion problem:

$$\begin{cases} \frac{\partial u}{\partial t}(x, t; y) - \nabla \cdot (\kappa(y)\nabla u(x, t; y)) + w(x, y) \cdot \nabla u(x, t; y) = 0, & (x, t) \in D \times (0, T), \\ u(x, t; y) = (1 - \exp(-t/0.1))(1 + 0.1 \sin(2\pi t)), & (x, t) \in \partial D \times (0, T), \\ u(x, 0; y) = 0, & x \in \bar{D}, \end{cases} \quad (40)$$

where the operator  $\nabla$  acts on the spatial variable  $x$  and the diffusion coefficient  $\kappa(y)$  and advection field  $w(x, y)$  are affine functions of the parameters  $y \in \Gamma$ , namely,

$$\kappa(y) := 0.1\kappa_0 + 1.8\kappa_0 y_1, \quad w(x, y) := w_0(x) + \sum_{i=1}^4 w_i(x)(2y_2 - 1), \quad (41)$$

with

$$w_0(x) := [2x_2(1 - x_1^2), -2x_2(1 - x_1^2)], \quad w_i(x) := \begin{cases} w_0(2(x - [c_i, d_i])) & x \in D_i, \\ [0, 0], & \text{otherwise,} \end{cases} \quad (42)$$

for  $c := [-0.5, -0.5, 0.5, 0.5]$ ,  $d := [-0.5, 0.5, -0.5, 0.5]$ ,  $D_i := [c_i - 0.5, d_i + 0.5] \times [c_i - 0.5, d_i + 0.5]$ . Three realisations of the advection field  $w(x, y)$  are given in Fig. 15.

Eq. (40) has a parametric diffusion coefficient and a parametric perturbation to the average advection field  $w_0$ . The solution does not reach a steady–state due to the time dependence in the Dirichlet boundary condition. Observe that for  $\kappa_0 > 0$  the problem is well posed because the diffusion coefficient is bounded  $0.09\kappa_0 < \kappa(y) < 2\kappa_0$  and the advection field is finite for all parameters  $y \in \Gamma$ , but never advection dominated. Further discussion on the well-posedness of this problem and verification of the analyticity of the solution  $u(x, t; y)$  with respect to the parameters  $y$  can be found in [81].

The quantity of interest is  $q(y) := u(x_*, t_*; y)$  for the space–time coordinate  $(x_*, t_*) := ([0.5, -0.5], 10)$ . Note that the chosen QoI is a pointwise quantity rather than an average or integrated quantity. This introduces additional difficulty as pointwise QoIs are typically more sensitive to local solution features and numerical errors, making accurate approximation more challenging.




Fig. 16. Realisations of the reference approximation  $u^{\text{ref}}(\mathbf{x}, 10; \mathbf{y})$  to the parametric parabolic PDE problem (40) at three parameter realisations  $\mathbf{y} = [-1, -1], [0, 0]$  and  $[1, 1]$ . The marker  $\times$  indicates the spatial location  $\mathbf{x}_* = (0.5, -0.5)$  where the QoI is evaluated.

4.1.1. Numerical approximation of PDE solution

Pointwise (in  $\Gamma$ ) approximations of the QoI  $q(\mathbf{y})$  are computed by using a  $P_1$  continuous Galerkin finite element (FE) approximation of (40) (piecewise linear polynomials on a structured uniform triangular mesh) and solving the resulting system of ODEs using TR-AB2 adaptive timestepping with local error control [82,83]. The solver accuracy (fidelity) is controlled by the local timestepping error tolerance  $\delta > 0$  whilst the mesh is fixed with  $400 \times 400$  uniformly spaced points. The maximum mesh Péclet number  $\mathcal{P}^h$  is obtained for the diffusion coefficient  $\kappa([0, y_2]) = 0.1\kappa_0$  and maximum advection field magnitude with  $y_2 = 1$ . This can be bound by  $\mathcal{P}_{\text{max}}^h = (h|\mathbf{w}_{\text{max}}|)/(2\kappa_{\text{min}}) = (4 \cdot 2 \cdot 400^{-1})/(2 \cdot 10^{-1} \cdot 10^{-1}) = 1$  and it has been numerically verified that the mesh provides stable solutions even in the worst case scenario. Cost is proportional to the number of timesteps used to reach  $t_* = 10$  which is approximately inversely proportional to the size of timesteps. For TR-AB2 the timesteps are approximately proportional to  $\delta^{1/3}$  so the solver cost is modelled as

$$\widehat{W}_\delta = \delta^{-1/3}. \tag{43}$$

There are two main sources of solver error: spatial and temporal. The spatial error in a semi-discrete solution using a conforming finite element approximation does not change the regularity of the computed approximation [7,84]. Conversely, the timestepping error does affect the regularity of the solution with respect to the parameter  $\mathbf{y}$ , and it is the main source of the spectral coefficient plateau in our computed approximation.

4.1.2. Multi-fidelity MISC approximation

The local timestepping error (LTE) tolerance  $\delta$  is used to define the different fidelities for the problem such that

$$\alpha \mapsto \{q^h : \Gamma \rightarrow \mathbb{R} \text{ computed with LTE tolerance } \delta = 10^{-\alpha}\} \tag{44}$$

with associated evaluation cost  $\widehat{W}_\alpha = 10^{\alpha/3}$  derived from (43). Symmetric Leja points and the two-step rule from (2) are used to form the MISC approximation.

As in Section 3.5, a reference approximation is constructed using an isotropic Smolyak sparse grid with multi-index set defined by (37) with  $w = 5$  and Clenshaw–Curtis points with the doubling rule from (2). The reference approximation uses the fidelity defined by  $\alpha^{\text{ref}} = 7$ . Three realisations of  $u^{\text{ref}}(\cdot, t_*; \mathbf{y})$  for parameters  $\mathbf{y} = [-1, -1], [0, 0]$  and  $[1, 1]$  are shown in Fig. 16.

The effects of changing the diffusion coefficient  $\kappa(\mathbf{y})$  and the advection field  $\mathbf{w}(\cdot, \mathbf{y})$  are clearly illustrated. For example, as the diffusion coefficient  $\kappa(\mathbf{y})$  increases, the diffusive boundary layer on the right hand wall becomes wider. Similarly, as the advection field  $\mathbf{w}(\cdot, \mathbf{y})$  increases in strength, the fluid is transported more strongly into the domain. Histograms of the approximation error of the QoI for  $\alpha = 1, 2, 3$  are shown in Fig. 17. It can be seen that as fidelity level  $\alpha$  increases, the approximation error decreases but the distribution is not necessarily Gaussian.

Computed response surfaces using the PLATEAUMISC and MISC algorithms are shown in Fig. 18 alongside pointwise error surfaces with respect to the reference approximation.

The pointwise error in the PLATEAUMISC approximation is generally more than one order of magnitude smaller than in the MISC approximation.

The benefits of PLATEAUMISC are shown more clearly in the error plots in Fig. 19. Note that the MISC and PLATEAUMISC approximations are not identical even for low computational cost: this is a consequence of constructing the output PLATEAUMISC approximation using only the multi-index set  $I$ , rather than the larger multi-index set  $I \cup R_I$ . The surfaces in Fig. 18 correspond to the last markers in Fig. 19. The estimated  $L^2_\rho(\Gamma; \mathbb{R})$  error is computed using (38) with  $N_{MC} = 10^4$  samples. For the adaptive




Fig. 17. Histograms of approximation error  $q^{\text{ref}} - q^\alpha$  from  $n = 1000$  Latin Hypercube samples for the parametric parabolic PDE (40) with fidelities  $q^\alpha$  defined in (44).

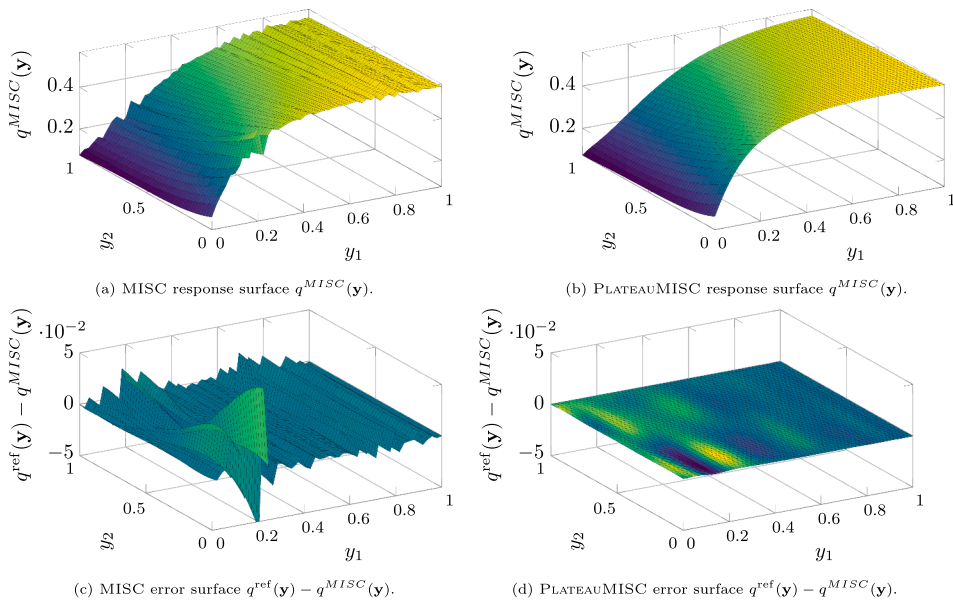


Fig. 18. Approximated response surfaces  $q^{\text{MISC}}(\mathbf{y})$  for the parametric parabolic PDE (40) using (a) Algorithm 1 and (b) Algorithm 3 with fidelities  $q^\alpha$  defined in (44). Fig. 18a and b illustrate the corresponding approximation error  $q^{\text{ref}}(\mathbf{y}) - q^{\text{MISC}}(\mathbf{y})$ .

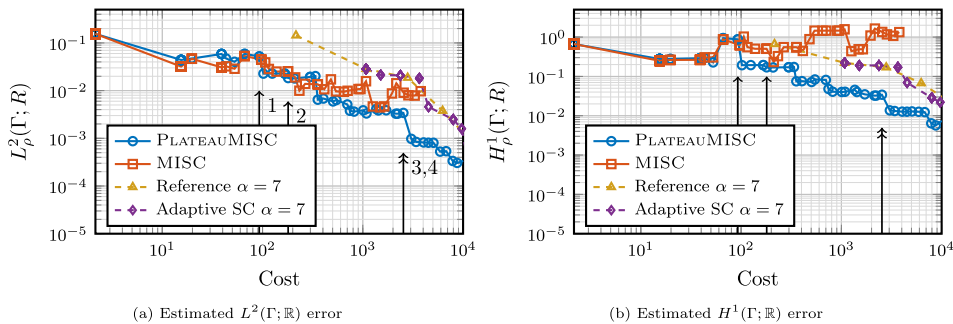
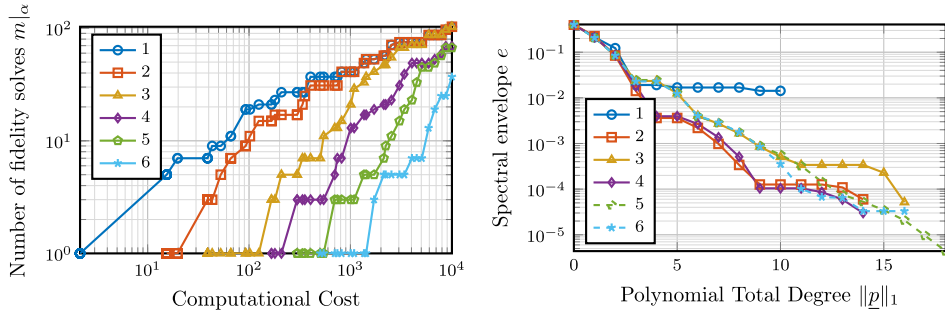


Fig. 19. Estimated  $L^2_\rho(\Gamma; \mathbb{R})$  and  $H^1_\rho(\Gamma; \mathbb{R})$  approximation error for the PLATEAUMISC approximation  $q^{\text{MISC}}$  of parametric parabolic PDE (40) using Algorithm 3 with fidelities  $q^\alpha$  defined in (44). Vertical arrows denote the times at which a fidelity is marked as saturated.



(a) Number of solves  $m_{|\alpha}$  for each fidelity  $\alpha$  at each iteration. (b) Solid lines denote spectral envelopes  $e_{|\alpha}$  for each fidelity  $\alpha$  at plateau detection (which happens at different iterations for different fidelities). Dashed lines denote spectral envelopes for  $\alpha = 5, 6$  at the final iteration, with no spectral plateau detected.

**Fig. 20.** Number of fidelity solves against computational work and snapshot of spectral polynomial coefficients for the PLATEAUMISC approximation  $q^{MISC}$  of the parametric parabolic PDE (40) using Algorithm 3 with fidelities  $q^\alpha$  defined in (44).

MISC algorithm, the  $L^2_\rho(\Gamma; \mathbb{R})$  error stagnates at approximately  $10^{-2}$  whilst the PLATEAUMISC approximation error decreases to  $3 \times 10^{-4}$ . Similarly, the  $H^1_\rho(\Gamma; \mathbb{R})$  error for the adaptive MISC algorithm stagnates at approximately  $10^0$  whilst the PLATEAUMISC approximation error decreases to  $10^{-2}$ . It is clear that the PLATEAUMISC algorithm gives a better approximation than the standard adaptive MISC approximation. The approximation errors are also compared to the sequential construction of the reference approximation with multi-index sets  $\{\beta \in \mathbb{N}^2 \text{ s.t. } \|\beta\|_1 \leq 2 + w\}$  for  $w = 0, 1, 2, 3, 4$  and fidelity level  $\alpha^{\text{ref}} = 7$ . The PLATEAUMISC approximation attains the same rate of convergence as the reference approximation, whilst achieving comparable approximation error for a reduced computational cost. This is because PLATEAUMISC can balance the solver error and parametric error giving a more efficiently constructed approximation. Similarly, using PLATEAUMISC we see an improvement over the adaptive single reference fidelity approximation.

The reduced computational cost is due to the balance between the solver error and parametric error. Initially, low fidelities can be used as the dominant error is the parametric interpolation error. In Fig. 20a we see that initially only fidelity  $\alpha = 1$  is sampled. As the total approximation error decreases, the dominant error becomes the solver error (i.e. time discretisation error). PLATEAUMISC detects this change, allowing higher fidelities to be accessed and reducing the overall approximation error further. The sequential activation of higher fidelities is shown in Fig. 20a. This is not the case for classic MISC as it is harder for higher fidelities to be activated, resulting in the stagnation of the error observed in Fig. 19.

The final detected set of saturated fidelities is  $S = \{1, 2, 3, 4\}$ . The spectral plateau effect is observed for fidelities  $\alpha = 1, 2, 3, 4$  in the spectral envelopes in Fig. 20b. The benefit of including the burn out parameter is also clear: in detecting the plateau for  $\alpha = 3$  it is clearly advantageous to drop the final two coefficients in the envelope. The spectral envelopes for fidelities  $\alpha = 5$  and  $\alpha = 6$  continue to decay, and the adaptive algorithm would continue to explore these fidelities.

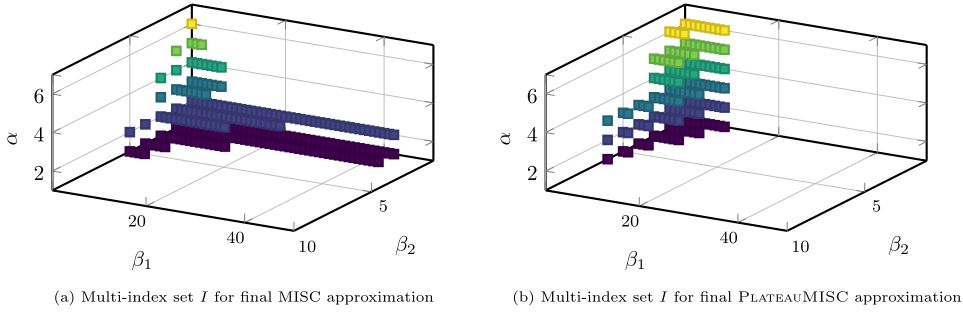
**Remark 2.** The spectral envelopes in Fig. 20b are represented at the iteration in which the plateau is first detected. After stagnation of a fidelity is detected, the algorithm can still add further collocation points to said fidelity, due to the back filling procedure using the multi-index set defined in (32). After adding additional collocation points, the corresponding spectral envelope may no longer show a plateau, but having been previously detected by PLATEAUMISC it remains marked as saturated. Finally, note that no spectral plateau is detected for fidelities  $\alpha = 5$  and  $\alpha = 6$ . Hence, the corresponding spectral envelopes (denoted by a dashed line in Fig. 20b) continue to decrease, with the associated solver noise not affecting the quality of the constructed surrogate model.

The multi-index sets  $I$  at the final algorithm iterations are shown in Fig. 21. Fig. 21a indicates that, in the classic MISC algorithm,  $\alpha = 1, 2$  are erroneously oversampled. The PLATEAUMISC algorithm prevents this, with the resulting multi-index set shown in Fig. 21b. The ‘sloped’ shape in the dimension  $\beta_2$  represents the lower fidelity information that is incorporated into the final approximation: higher degree polynomial terms are used in fidelities  $\alpha = 1, 2, 3$  than  $\alpha = 4, 5, 6$ .

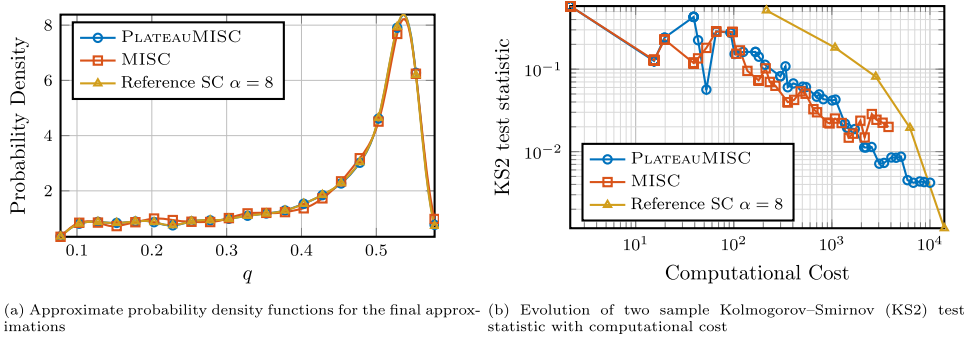
The PDFs are again approximated using  $N_{MC} = 10^4$  samples with the MATLAB `ksdensity` function and the KS2 test statistic is computed with  $N_{MC} = 10^6$  samples. In Fig. 22, the distinction between MISC and PLATEAUMISC PDFs is less clear than in the analytic example of Section 3.5, but PLATEAUMISC still outperforms classic MISC in terms of accuracy, efficiency and robustness.

#### 4.2. Parametric flow control in turbulent incompressible Navier–Stokes problem

This example considers the NASA two dimensional wall mounted hump (2DWMH) benchmark test case [78,79,85], which simulates a turbulent incompressible flow over a Glauert–Goldschmied-type body controlled by a parameter  $y \in \Gamma = [0, 1]^2$ . The problem



**Fig. 21.** Multi-index sets  $I$  for the final MISC and PLATEAUMISC approximations of the parametric parabolic PDE (40) using Algorithms 1 and 3 with fidelities  $q^\alpha$  defined in (44).



**Fig. 22.** Final approximate probability density functions and evolution of two sample Kolmogorov–Smirnov test statistic for the parametric parabolic PDE (40) using Algorithm 3 with fidelities  $q^\alpha$  defined in (44).

is governed by the Reynolds-averaged Navier–Stokes equations for incompressible flows, with Spalart–Allmaras turbulence model:

$$\begin{cases} \frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot (\mathbf{u} \otimes \mathbf{u}) - \nabla \cdot (2(\nu + \nu_t) \nabla^s \mathbf{u}) + \nabla p = \mathbf{0}, & (\mathbf{x}, t) \in D \times (0, T), \\ \nabla \cdot \mathbf{u} = 0, & (\mathbf{x}, t) \in D \times (0, T), \\ \frac{\partial \tilde{\nu}}{\partial t} + \nabla \cdot (\mathbf{u} \tilde{\nu}) - \frac{1}{\sigma} \nabla \cdot ((\nu + \tilde{\nu}) \nabla \tilde{\nu}) - \frac{c_{b2}}{\sigma} \nabla \tilde{\nu} \cdot \nabla \tilde{\nu} - c_{b1} \tilde{\nu} \tilde{\nu} + c_{w1} f_w \left( \frac{\tilde{\nu}}{d} \right)^2 = 0, & (\mathbf{x}, t) \in D \times (0, T), \end{cases} \quad (45)$$

where  $\mathbf{u}(\mathbf{x}, t; \mathbf{y})$ ,  $p(\mathbf{x}, t; \mathbf{y})$  and  $\tilde{\nu}(\mathbf{x}, t; \mathbf{y})$  are the parameter-dependent velocity, pressure and eddy viscosity fields. Moreover, the turbulent viscosity is defined as

$$\nu_t = \tilde{\nu} f_{v1}, \quad (46)$$

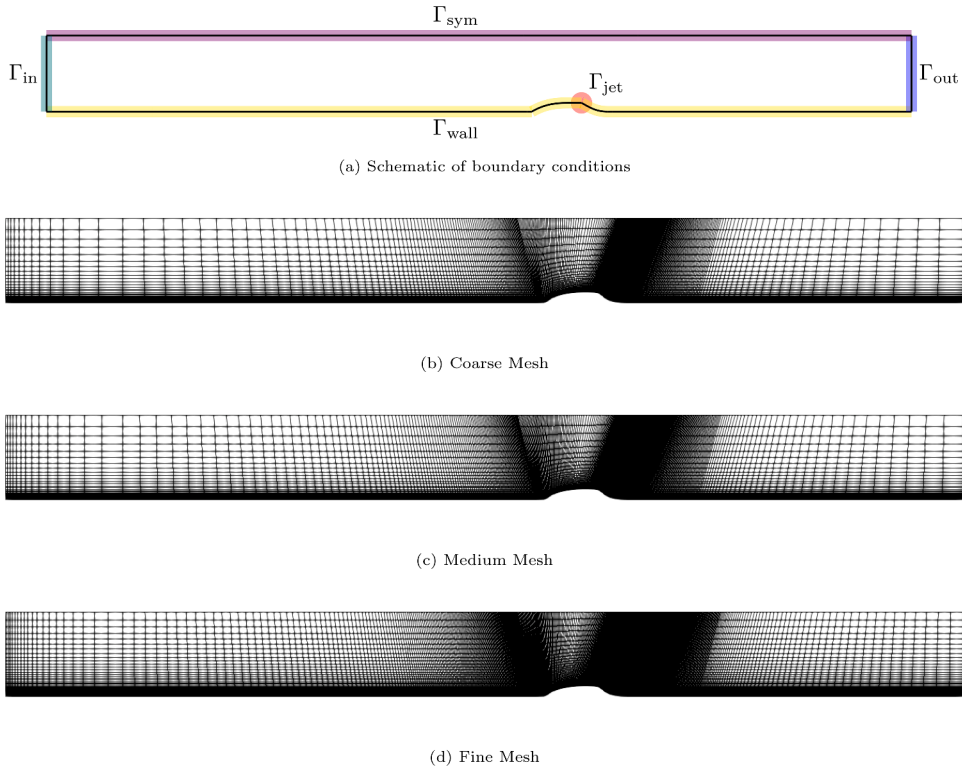
and  $\nu$ ,  $\nabla^s \mathbf{u} = (\nabla \mathbf{u} + (\nabla \mathbf{u})^\top)/2$  and  $\tilde{d}$  respectively denote the physical viscosity of the fluid, the strain-rate tensor and the distance from the closest wall in the domain. Finally, the Spalart–Allmaras model is completed with a set of closure functions [86]:

$$\begin{aligned} \Omega &:= \frac{\nabla \mathbf{u} - (\nabla \mathbf{u})^\top}{2}, & \tilde{S} &:= [2\Omega : \Omega]^{1/2} + \frac{\tilde{\nu}}{\kappa^2 \tilde{d}^2} f_{v2}, & \chi &:= \frac{\tilde{\nu}}{\nu}, \\ f_w &:= g \left[ \frac{1 + c_{w3}^6}{g^6 + c_{w3}^6} \right]^{1/6}, & f_{v2} &:= 1 - \frac{\chi}{1 + \chi f_{v1}}, & f_{v1} &:= \frac{\chi^3}{\chi^3 + c_{v1}^3}, \\ g &:= r + c_{w2}(r^6 - r), & r &:= \frac{\tilde{\nu}}{\tilde{S} \kappa^2 \tilde{d}^2}, \end{aligned} \quad (47)$$

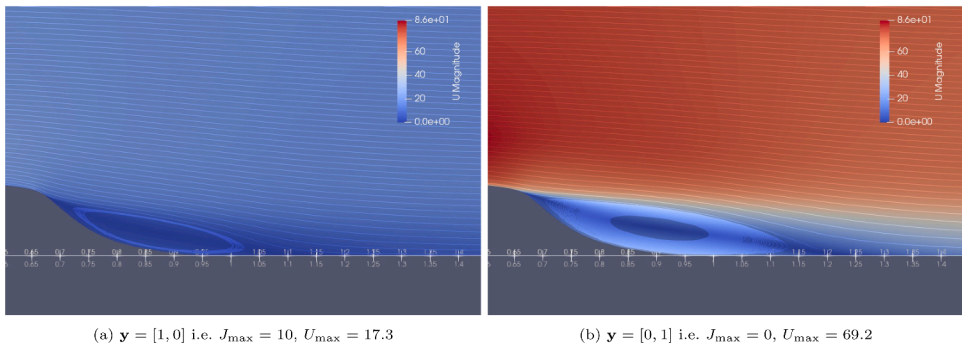
where  $\sigma = 2/3$ ,  $c_{b1} := 0.1355$ ,  $c_{b2} = 0.622$ ,  $\kappa = 0.41$ ,  $c_{w1} = c_{b1}/\kappa^2 + (1 + c_{b2})/\sigma$ ,  $c_{v1} = 7.1$ ,  $c_{w2} = 0.34$  and  $c_{w3} = 2$ .

The geometry of the spatial domain  $D$ , the set of boundary conditions and three example meshes are shown in Fig. 23. Let  $L = 0.42m$  denote the chord length of the hump, with its maximum height set to  $H = 0.0535m$ . The computational domain  $D$  is a channel of height  $L$  extending  $6.39L$  upstream and  $5L$  downstream of the hump. At the inlet  $\Gamma_{in}$ , a parabolic profile is imposed on both velocity and turbulent viscosity, with the maximum velocity magnitude parameterised by  $U_{max} : \Gamma \rightarrow \mathbb{R}$ ,

$$U_{max}(\mathbf{y}) := 17.3 + (69.2 - 17.3)y_2, \quad (48)$$



**Fig. 23.** Geometry of spatial domain  $D$ , boundary conditions and three examples of spatial meshes for the NASA 2DWMH test problem described in Section 4.2.



**Fig. 24.** Velocity streamlines for two parameter realisations  $\mathbf{y} = [1, 0]$  and  $\mathbf{y} = [0, 1]$  in the region behind the hump.

whereas the eddy viscosity is set to  $\tilde{\nu} = 3\nu$ , assuming fully developed turbulence. A suction jet is located at approximately 65% of the chord length  $L$  on a patch named  $\Gamma_{\text{jet}}$ . On  $\Gamma_{\text{jet}}$ , homogeneous Neumann conditions are imposed on the eddy viscosity and a smooth velocity profile is defined, see [80], with maximum outflow magnitude  $J_{\max} : \Gamma \rightarrow \mathbb{R}$ ,

$$J_{\max}(\mathbf{y}) := 10y_1. \tag{49}$$


Finally, homogeneous Dirichlet conditions are enforced for both  $\mathbf{u}$  and  $\tilde{\nu}$  on  $\Gamma_{\text{wall}}$ , homogeneous Neumann conditions are applied on  $\Gamma_{\text{out}}$  and a symmetry condition is applied on  $\Gamma_{\text{sym}}$ .

The kinematic viscosity of the fluid is set to  $\nu = 1.55274 \cdot 10^{-5} \text{ m}^2/\text{s}$ . Setting the chord length  $L$  as the characteristic length of the problem and the characteristic velocity equal to the maximum inlet velocity, the resulting Reynolds number  $\text{Re} := (U_{\max} L)/\nu$  ranges between  $4.70 \cdot 10^5$  and  $1.88 \cdot 10^6$  — which is clearly in the turbulent regime.

The parametric formulation is motivated by a desire to control the flow reattachment point (i.e., the QoI  $q : \Gamma \rightarrow \mathbb{R}$ ) by using a suction jet of varying magnitude at the back of the hump. Example streamlines for two flow realisations are shown in Fig. 24, demonstrating the variability in the flow separation and the flow reattachment point.

**Table 1**  
OpenFOAM `blockMesh` configuration parameters for the coarse, medium and fine meshes in the 2DWMH problem in Section 4.2.

Fidelity	Coarse	Medium	Fine
$\alpha$	1	2	3
Channel height ( $L$ )	1	1	1
Number of cells in vertical direction	42	48	96
Height of first cell	$4.25 \times 10^{-4}$	$1.83 \times 10^{-4}$	$1.00 \times 10^{-5}$
Height of last cell	$1.31 \times 10^{-1}$	$1.31 \times 10^{-1}$	$9.16 \times 10^{-2}$
Total expansion ratio (last height / first height)	308	712	9156
Cell-to-cell expansion ratio	1.15	1.15	1.10
Range of $y^+$ for $U_{\max} = 69.2, J_{\max} = 0$	[6.96, 23.78]	[2.41, 13.12]	[0.09, 0.99]



**Fig. 25.** Computed  $y^+$  values for the coarse, medium and fine meshes in the 2DWMH problem in Section 4.2. The coordinate  $x/L$  is the scaled  $x$  coordinate along the bottom boundary where  $L$  is the hump length and the origin is the start of the hump.

4.2.1. Numerical approximation of the flow fields

To compute the reattachment point, (45) is solved using a cell-centered finite volume scheme with pseudo-transient relaxation, and a staggered algorithm splitting the flow equations and turbulence model. The `simpleFoam` solver in OpenFOAM [87] provides an implementation of the SIMPLE algorithm [88] to simulate such a problem. Note that, although the problem under analysis is steady-state, the `simpleFoam` solver employs a pseudo-transient formulation to treat the non-linear terms, thus motivating the presence of the time derivative in (45).

The reattachment point is postprocessed by evaluating the wall shear stress along the lower boundary: the point where the shear stress changes sign indicates reattachment. This computation requires accurate velocity fields, especially near the wall. Indeed, it is well known that a value of the dimensionless wall distance  $y^+ < 1$  is required in order to capture the physical phenomena in the vicinity of the walls. In the following subsection, it will be shown that the PLATEAUMISC algorithm can extract relevant information from coarse meshes ( $y^+ \gg 1$ ), normally unsuitable for turbulent simulations, automatically identifying when the potential of a mesh has been exhausted. These under-resolved meshes are expected to provide quantitatively inaccurate, and possibly noisy, predictions of the quantity of interest.

4.2.2. Multi-fidelity MISC approximation

Three fidelities are defined using the coarse, medium and fine meshes (see Fig. 23), each with distinct horizontal and vertical cell grading, and solver tolerances. The vertical `blockMesh` grading details are listed in Table 1.

The mesh size at the wall and the expansion ratio play a critical role in the accuracy of the solution and, consequently, in the computation of the reattachment point. As previously mentioned, a value of the dimensionless height of the first boundary cell,  $y^+$ , below 1 is required to fully resolve the boundary layer without using wall functions. Fig. 25 shows that this is only verified for the fine mesh.


Hence, one would expect that only the fine mesh can accurately compute the reattachment point. Nonetheless, under-resolved meshes can be employed by PLATEAUMISC to extract some useful information from the problem. The solver tolerances are listed in Table 2.

A combination of the mesh size and iterative solver tolerances controls the solver error. In Fig. 26 approximation error histograms for the fidelities  $\alpha = 1, 2$  are shown, with errors computed relative to the finest fidelity  $\alpha = 3$ . This shows that the solver error decreases as the fidelity  $\alpha$  is increased, although no clear distribution emerges. Note also that the histograms are not centered around zero, i.e. both fidelities have a non-negligible bias.

To quantify computational cost for each fidelity, we average the solver runtime over a sample of parameter realisations: specifically samples from a sparse grid of Leja points with the two-step rule from (2) and isotropic multi-index set  $\{\beta \in \mathbb{N}^2 \text{ s.t. } \|\beta\|_1 \leq 2 + w\}$  with  $w = 6$ . All simulations were performed on a workstation with an Intel Xeon Silver 4215R CPU @ 3.20GHz. Table 2 presents the estimated costs: the lowest fidelity is roughly five times faster than the intermediate fidelity, and 137 times faster than the finest fidelity.

**Table 2**  
OpenFOAM simpleFoam configuration parameters for the coarse, medium and fine meshes in the 2DWMH problem in Section 4.2.

Fidelity	Coarse	Medium	Fine
$\alpha$	1	2	3
Cells	18,060	24,384	69,120
Tolerances (residual)			
velocity	$10^{-3}$	$10^{-4}$	$10^{-6}$
turbulence	$10^{-3}$	$10^{-4}$	$10^{-6}$
pressure	$10^{-2}$	$10^{-3}$	$5 \times 10^{-6}$
Tolerances (inner loop solver)			
absolute	$10^{-6}$	$10^{-6}$	$10^{-10}$
relative	$10^{-2}$	$10^{-2}$	$10^{-4}$
Computational cost (sample mean solve time)	36.9 s	183.1 s	5062.8 s



**Fig. 26.** Histograms of approximation error  $q^{\text{ref}} - q^\alpha$  from  $n = 1000$  Latin Hypercube samples for the 2DWMH problem using fidelities  $q^\alpha$  defined in Table 2.

The QoI (the reattachment point) is approximated using MISC and PLATEAUMISC, see Algorithms 1, 3 respectively. A reference approximation is computed with isotropic Smolyak multi-index set (37) with  $w = 5$ , Clenshaw–Curtis points and the doubling rule, and the fine fidelity  $\alpha = 3$  defined in Table 2. The final response surfaces are shown in Fig. 27. The PLATEAUMISC approximation is clearly more accurate and less corrupted by solver error than the classical MISC strategy.

The  $L^2_\rho(\Gamma; R)$ , and similarly the  $H^1_\rho(\Gamma; \mathbb{R})$ , norms of the approximation error are estimated numerically via (38) with  $N_{MC} = 10^4$  Monte Carlo samples. These are presented in Fig. 28.

At low computational budgets, both approximations are poor but are informative of underlying trends and approximate values; note that at such a low cost, no approximation is attainable using only the high fidelity solver. As the computational budget increases, the PLATEAUMISC approximation error decreases at a rate comparable to the reference approximation, whereas the MISC approximation error stagnates and does not give a useful approximation, with large errors even at large computational budgets. This again demonstrates that PLATEAUMISC is robust to the solver errors that are hindering the accuracy of the classical MISC approximation. Algorithm 3 is able to balance solver cost and approximation error to automatically increase the fidelity and extract as much information as possible for a given computational budget.

Sampling profiles and spectral envelopes for the PLATEAUMISC approximation are shown in Fig. 29.

Fig. 29a in combination with Fig. 28 illustrates that substantial error reduction only occurs when the number of samples from the high-fidelity ( $\alpha = 3$ ) increases. Spectral envelopes in Fig. 29b (despite not having theoretical proof that the QoI is analytic with respect to the parameters) exhibit an exponential decay followed by spectral plateaus. This allows PLATEAUMISC to successfully identify stagnation and avoid overfitting to the solver noise, unlike classical MISC.

It is worth highlighting that PLATEAUMISC is capable of using few samples from under-resolved meshes with  $y^+ = O(10^1)$  to extract relevant features of the flow, without “overfitting” the solver error. This allows us to improve the accuracy of the multi-fidelity surrogate (see Fig. 27), and also significantly reduces the evaluations of all configurations. The number of evaluations of fidelity  $\alpha = 3$ , the only fidelity with  $y^+ < 1$ , is reduced from 41 evaluations in classic MISC to 27 evaluations in PLATEAUMISC. The final multi-index sets, shown in Fig. 30, also illustrate the tendency to oversample the lower fidelities by the classic MISC algorithm.

Finally, Fig. 31a compares estimates of the probability density function of the QoI, constructed using  $N_{MC} = 10^4$  Monte Carlo samples. The MISC approximation deviates significantly from the reference, while PLATEAUMISC matches it closely. Fig. 31b shows the evolution of the KS2 test statistic which stagnates for the MISC approximation, but decreases consistently for PLATEAUMISC illustrating its superior performance.

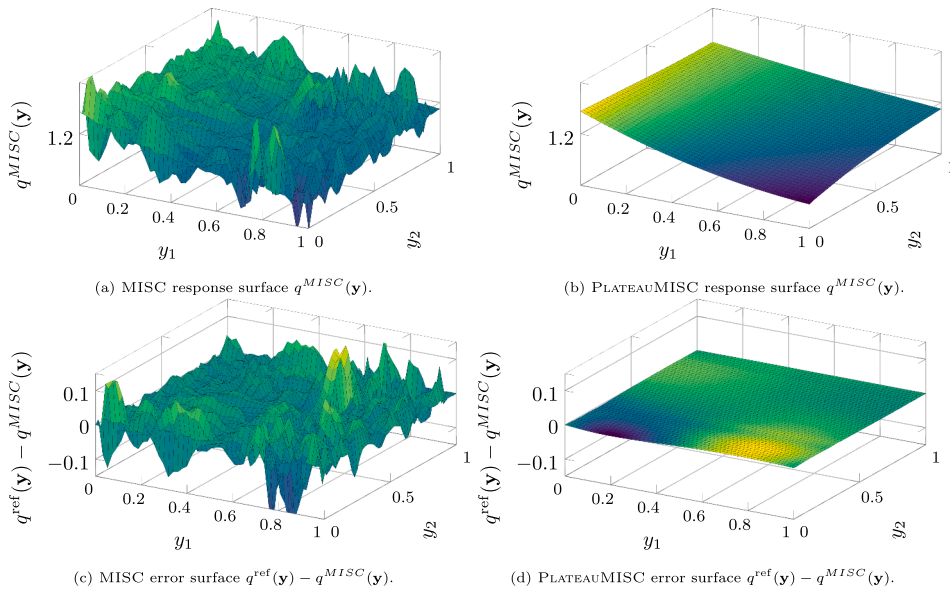


Fig. 27. Approximated response surfaces  $q^{MISC}(\mathbf{y})$  for the 2DWMH problem using (a) Algorithm 1 and (b) Algorithm 3 with fidelities  $q^\alpha$  defined in Table 2. Fig. 27c and d illustrate the corresponding approximation errors  $q^{ref}(\mathbf{y}) - q^{MISC}(\mathbf{y})$ .

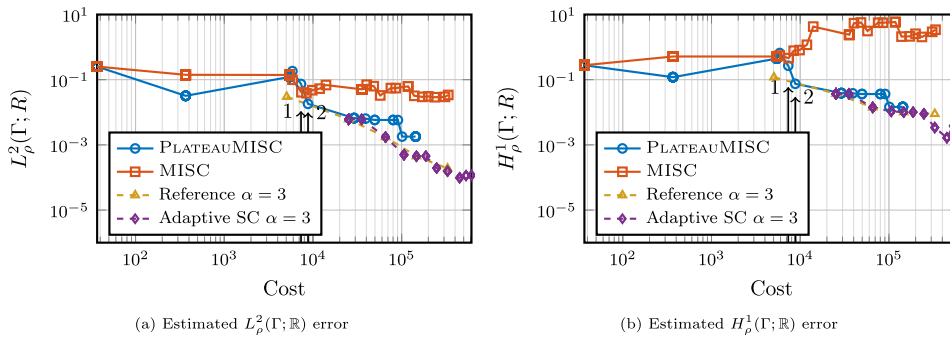


Fig. 28. Estimated  $L^2_p(\Gamma; \mathbb{R})$  and  $H^1_p(\Gamma; \mathbb{R})$  approximation error for the PLATEAUMISC approximation  $q^{MISC}$  for the 2DWMH problem using Algorithm 3 with fidelities  $q^\alpha$  defined in Table 2. Vertical arrows denote the times at which a fidelity is marked as saturated.

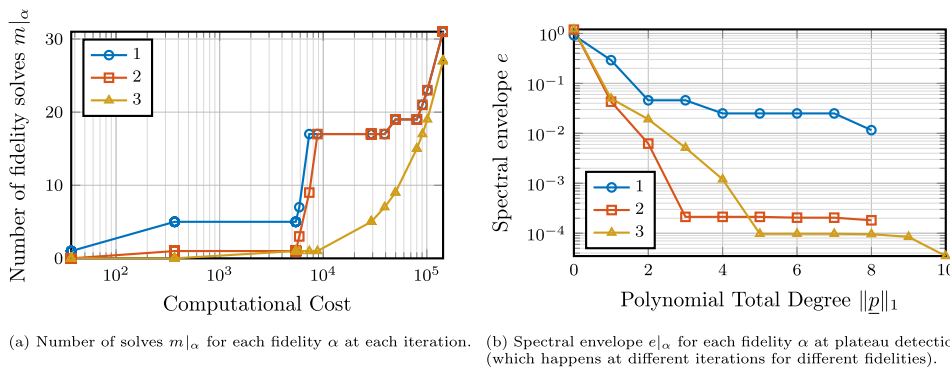
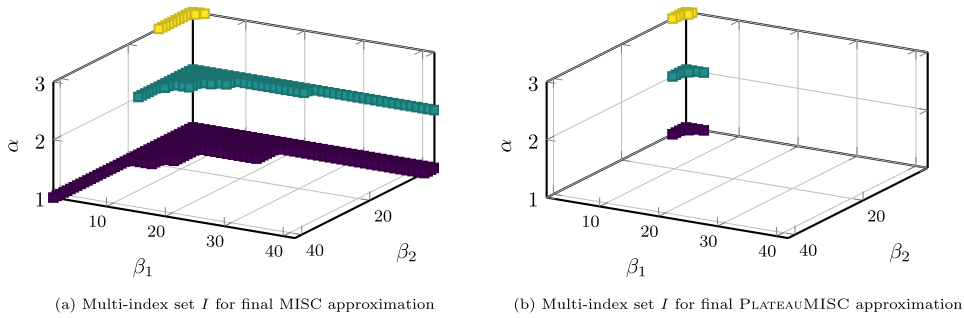
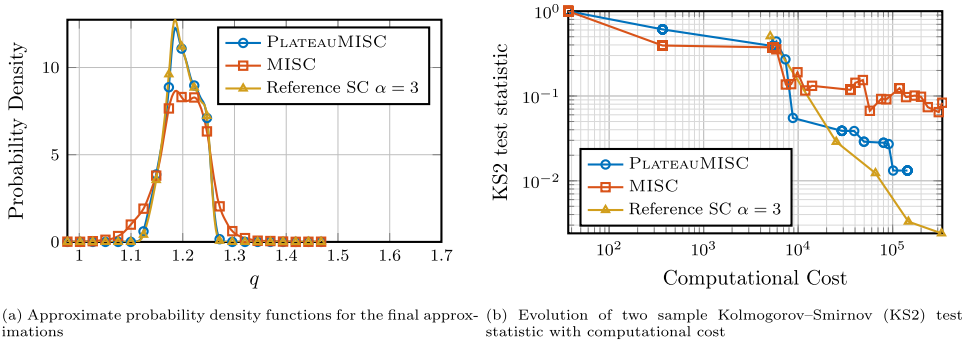


Fig. 29. Number of fidelity solves against computational work and snapshot of spectral polynomial coefficients for the PLATEAUMISC approximation  $q^{MISC}$  of the 2DWMH problem using Algorithm 3 with fidelities  $q^\alpha$  defined in Table 2.



**Fig. 30.** Multi-index sets  $I$  for the MISC and PLATEAUMISC approximation  $q^{MISC}$  of the 2DWMH prolem using Algorithms 1 and 3 with fidelities  $q^a$  defined in Table 2.



**Fig. 31.** Final approximate probability density functions and evolution of two sample Kolmogorov–Smirnov test statistic for the PLATEAUMISC approximation  $q^{MISC}$  of the 2DWMH prolem using Algorithm 3 with fidelities  $q^a$  defined in Table 2.

### 5. Conclusions

Using the standard adaptive multi-index stochastic collocation algorithm to compute surrogates of parametric PDEs can result in an approximation whose quality stagnates since the algorithm can not distinguish between the solver error and the true parametric response that must be captured. The PLATEAUMISC algorithm can identify the effect of solver noise by inspecting the spectral content of the constructed MISC approximation: a plateau detection algorithm identifies and blocks further refinement in the saturated fidelities, resulting in a multi-fidelity MISC approximation robust to such intrinsic noise.

The accuracy of the PLATEAUMISC algorithm is demonstrated through an analytic Gaussian peak example, a parabolic PDE example with uncertain coefficients and a parametric turbulent Navier–Stokes flow. In all of these examples the PLATEAUMISC algorithm is able to converge at a rate comparable to the high fidelity reference approximation. In the analytic Gaussian example and the parametric parabolic PDE example, a significant reduction in computational cost is observed when compared to a single high fidelity approximation. In the parametric turbulent flow example, PLATEAUMISC numerically demonstrates the ability to automatically select the appropriate fidelities required to construct a robust surrogate model for quantities of engineering interest, whilst extracting relevant information even from under-resolved meshes not suitable for reliable single fidelity computations. In other words, the advantage of multi-fidelity approximation over single-fidelity is dependent on the pool of available fidelities and PLATEAUMISC is able to exploit whatever advantage is there due to its robustness to noise, whilst classic MISC may result in suboptimal approximations overfitting intrinsic solver errors.

Finally, we remark that this work has proposed a robust variant of MISC, outperforming the standard one in the presence of noisy evaluations of quantities of interests. Whilst we can point out that the reduced number of hyperparameters to be tuned represents an appealing feature of PLATEAUMISC also with respect to other approaches (e.g., based on least-squares techniques), a thorough benchmarking is required to globally assess advantages and disadvantages of the different noise-robust multi-fidelity strategies in the literature. In this context, a benchmarking campaign in the spirit of [89,90] is expected to significantly contribute to a deeper understanding of the performance of noise-robust multi-fidelity methods, including PLATEAUMISC. For this purpose, a crucial point is represented by the selection of appropriate test problems accounting for the different aspects possibly influencing each method, including, e.g., choice of samples, choice of polynomial spaces, and algorithms to allocate over fidelities, just to name a few.

### CRedit authorship contribution statement

**Benjamin M. Kent:** Writing – review & editing, Writing - original draft, Software, Methodology, Investigation, Formal analysis, Conceptualization; **Lorenzo Tamellini:** Writing – review & editing, Supervision, Software, Methodology, Investigation, Funding

acquisition, Formal analysis, Conceptualization; **Matteo Giacomini**: Writing – review & editing, Supervision, Software, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization; **Antonio Huerta**: Writing – review & editing, Supervision, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

## Data availability

Data will be made available on request.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

B. Kent and L. Tamellini have been partially supported by the project 202222PACR “Numerical approximation of uncertainty quantification problems for PDEs by multi-fidelity methods (UQ-FLY)”, funded by European Union–NextGenerationEU. L. Tamellini has been partially supported by the project PRIN PNRR “Uncertainty Quantification of coupled models for water flow and contaminant transport” (P2022LXLYY), financed by the European Union–NextGeneration EU. L. Tamellini is member of the Gruppo Nazionale Calcolo Scientifico-Istituto Nazionale di Alta Matematica (GNCS-INdAM).

M. Giacomini and A. Huerta acknowledge the support of the Spanish Ministry of Science, Innovation and Universities and Spanish State Research Agency MICIU/AEI/10.13039/501100011033 (Grant No. PID2023-149979OB-I00) and of the Generalitat de Catalunya (Grant No. 2021- SGR-01049). M. Giacomini is Fellow of the Serra Hünter Programme of the Generalitat de Catalunya.

## References

- [1] R. Ghanem, D. Higdon, H. Owhadi, Handbook of Uncertainty Quantification, Handbook of Uncertainty Quantification, Springer International Publishing, 2017. <https://doi.org/10.1007/978-3-319-12385-1>
- [2] R.C. Smith, Uncertainty Quantification: Theory, Implementation, and Applications, Computational Science and Engineering, SIAM, 2013. <https://doi.org/10.1137/1.9781611973228>
- [3] T.J. Sullivan, Introduction to Uncertainty Quantification, Texts in Applied Mathematics, Springer, Cham, Switzerland, 1 edition, Cham, Switzerland, 2015. <https://doi.org/10.1007/978-3-319-23395-6>
- [4] R.B. Gramacy, Surrogates: Gaussian Process Modeling, Design and Optimization for the Applied Sciences, Chapman Hall/CRC, Boca Raton, Florida, Boca Raton, Florida, 2020. <https://doi.org/10.1201/9780367815493>
- [5] J.S. Hesthaven, G. Rozza, B. Stamm, Certified Reduced Basis Methods for Parametrized Partial Differential Equations, Springer Cham, 2015. <https://doi.org/10.1007/978-3-319-22470-1>
- [6] A. Quarteroni, A. Manzoni, F. Negri, Reduced Basis Methods for Partial Differential Equations: an Introduction, Springer, 2015. <https://doi.org/10.1007/978-3-319-15431-2>
- [7] B. Adcock, S. Brugiapaglia, C.G. Webster, Sparse Polynomial Approximation of High-Dimensional Functions, Society for Industrial and Applied Mathematics, Philadelphia, PA, Philadelphia, PA, 2022. <https://doi.org/10.1137/1.9781611976885>
- [8] G. Kutyniok, P. Petersen, M. Raslan, R. Schneider, A theoretical analysis of deep neural networks and parametric PDEs, Construct. Approximation (2022). <https://doi.org/10.1007/s00365-021-09551-4>
- [9] J.S. Hesthaven, S. Ubbiali, Non-intrusive reduced order modeling of nonlinear problems using neural networks, J. Comput. Phys. 363 (2018) 55–78. <https://doi.org/10.1016/j.jcp.2018.02.037>
- [10] S. Fresca, L. Dedé, A. Manzoni, A comprehensive deep learning-Based approach to reduced order modeling of nonlinear time-dependent parametrized PDEs, J. Sci. Comput. 87 (2021) 61. <https://doi.org/10.1007/s10915-021-01462-7>
- [11] B. Adcock, N. Dexter, The gap between theory and practice in function approximation with deep neural networks, SIAM J. Math. Data Sci. 3 (2) (2021) 624–655. <https://doi.org/10.1137/20M131309X>
- [12] F. Chinesta, A. Huerta, G. Rozza, K. Willcox, Model Reduction Methods, Wiley, 2017, p. 1–36. <https://doi.org/10.1002/9781119176817.ecm2110>
- [13] M. Giacomini, L. Borchini, R. Sevilla, A. Huerta, Separated response surfaces for flows in parametrised domains: comparison of a priori and a posteriori PGD algorithms, Finite Elem. Anal. Des. 196 (2021) 103530. <https://doi.org/10.1016/j.finel.2021.103530>
- [14] A. Chkifa, A. Cohen, C. Schwab, et al., High-dimensional adaptive sparse polynomial interpolation and applications to parametric PDEs, Found. Comput. Math. 14 (4) (2014) 601–633. <https://doi.org/10.1007/s10208-013-9154-z>
- [15] M. Eigel, C.J. Gittelsohn, C. Schwab, E. Zander, Adaptive stochastic Galerkin FEM, Comput. Methods Appl. Mech. Eng. 270 (2014) 247–269. <https://doi.org/10.1016/j.cma.2013.11.015>
- [16] F. Nobile, L. Tamellini, F. Tesei, R. Tempone, An adaptive sparse grid algorithm for elliptic PDEs with lognormal diffusion coefficient, in: J. Garcke, D. Pflüger (Eds.), Sparse Grids and Applications - Stuttgart 2014, Springer International Publishing, Cham, Switzerland, 2016, pp. 191–220. <https://doi.org/10.1007/978-3-319-28262-6>
- [17] A. Bespalov, D. Silvester, Efficient adaptive stochastic Galerkin methods for parametric operator equations, SIAM J. Sci. Comput. 38 (4) (2016) A2118–A2140. <https://doi.org/10.1137/15m1027048>
- [18] J.D. Jakeman, M.S. Eldred, G. Geraci, A. Gorodetsky, Adaptive multi-index collocation for uncertainty quantification and sensitivity analysis, Int. J. Numer. Methods Eng. 121 (6) (2019) 1314–1343. <https://doi.org/10.1002/nme.6268>
- [19] A.-L. Haji-Ali, F. Nobile, L. Tamellini, R. Tempone, Multi-index stochastic collocation convergence rates for random PDEs with parametric regularity, Found. Comput. Math. 16 (6) (2016) 1555–1605. <https://doi.org/10.1007/s10208-016-9327-7>
- [20] A.J. Crowder, C.E. Powell, A. Bespalov, Efficient adaptive multilevel stochastic Galerkin approximation using implicit a posteriori error estimation, SIAM J. Sci. Comput. 41 (3) (2019) A1681–A1705. <https://doi.org/10.1137/18m1194420>
- [21] A. Bespalov, D.J. Silvester, Error estimation and adaptivity for stochastic collocation finite elements part II: multilevel approximation, SIAM J. Sci. Comput. 45 (2) (2023) A781–A797. <https://doi.org/10.1137/22m1479361>
- [22] L.L. Gratiet, J. Garnier, Recursive co-kriging model for design of computer experiments with multiple levels of fidelity, Int. J. Uncertain Quantif. 4 (5) (2014) 365–386. <https://doi.org/10.1615/Int.J.UncertaintyQuantification.2014006914>
- [23] L.W.-T. Ng, M. Eldred, Multifidelity Uncertainty Quantification Using Non-Intrusive Polynomial Chaos and Stochastic Collocation, American Institute of Aeronautics and Astronautics, 2012. <https://doi.org/10.2514/6.2012-1852>

- [24] B. Peherstorfer, K. Willcox, M. Gunzburger, Survey of multifidelity methods in uncertainty propagation, inference, and optimization, *SIAM Rev.* 60 (3) (2018) 550–591. <https://doi.org/10.1137/16M1082469>
- [25] A.-L. Haji-Ali, F. Nobile, L. Tamellini, R. Tempone, Multi-index stochastic collocation for random PDEs, *Comput. Methods Appl. Mech. Eng.* 306 (2016) 95–122. <https://doi.org/10.1016/j.cma.2016.03.029>
- [26] I. Babuška, F. Nobile, R. Tempone, A stochastic collocation method for elliptic partial differential equations with random input data, *SIAM Rev.* 52 (2) (2010) 317–355. <https://doi.org/10.1137/100786356>
- [27] D. Xiu, J.S. Hesthaven, High-order collocation methods for differential equations with random inputs, *SIAM J. Sci. Comput.* 27 (3) (2005) 1118–1139. <https://doi.org/10.1137/040615201>
- [28] T. Gerstner, M. Griebel, Dimension-adaptive tensor-product quadrature, *Computing* 71 (1) (2003) 65–87. <https://doi.org/10.1007/s00607-003-0015-5>
- [29] C. Piazzola, L. Tamellini, Algorithm 1040: the sparse grids MATLAB kit - a MATLAB implementation of sparse grids for high-dimensional function approximation and uncertainty quantification, *ACM Trans. Math. Software* 50 (1) (2024). <https://doi.org/10.1145/3630023>
- [30] C. Piazzola, L. Tamellini, R. Pellegrini, R. Broglia, A. Serani, M. Diez, Comparing multi-index stochastic collocation and multi-fidelity stochastic radial basis functions for forward uncertainty quantification of ship resistance, *Eng. Comput.* (2022). <https://doi.org/10.1007/s00366-021-01588-0>
- [31] M. Chiappetta, C. Piazzola, L. Tamellini, A. Reali, F. Auricchio, M. Carraturo, Data-informed uncertainty quantification for laser-based powder bed fusion additive manufacturing, *Int. J. Numer. Methods Eng.* 125 (17) (2024). <https://doi.org/10.1002/nme.7542>
- [32] S. Brown, J. Beck, H. Mahgerefteh, E.S. Fraga, Global sensitivity analysis of the impact of impurities on CO2 pipeline failure, *Reliabil. Eng. Syst. Saf.* 115 (2013) 43–54. <https://doi.org/10.1016/j.res.2013.02.006>
- [33] J.H.S. de Baar, B. Harding, A gradient-enhanced sparse grid algorithm for uncertainty quantification, *Int. J. Uncertain Quantif.* 5 (5) (2015) 453–468. <https://doi.org/10.1615/int.j.uncertaintyquantification.2015014394>
- [34] R. Pellegrini, J. Wackers, R. Broglia, A. Serani, M. Visonneau, M. Diez, A multi-fidelity active learning method for global design optimization problems with noisy evaluations, *Eng. Comput.* (2023) 3183–3206. <https://doi.org/10.1007/s00366-022-01728-0>
- [35] T. Lukaczyk, T. Taylor, F. Palacios, J. Alonso, Managing Gradient Inaccuracies while Enhancing Optimal Shape Design Methods, *American Institute of Aeronautics and Astronautics*, 2013. <https://doi.org/10.2514/6.2013-1042>
- [36] B.W. Reuter, G. Geraci, T. Wildey, Analysis of the challenges in developing sample-based multifidelity estimators for nondeterministic models, *Int. J. Uncertain Quantif.* 14 (5) (2024) 1–30. <https://doi.org/10.1615/Int.J.UncertaintyQuantification.2024050125>
- [37] Haji-Ali Abdul-Lateef, Nobile Fabio, Tempone Raúl, Wolfers Sören, Multilevel weighted least squares polynomial approximation. *ESAIM: M2AN* 2020;54(2):649–677. <https://doi.org/10.1051/m2an/2019045>.
- [38] G. Migliorati, F. Nobile, E. von Schwerin, R. Tempone, Analysis of discrete  $L^2$  projection on polynomial spaces with random evaluations, *Found. Comput. Math.* 14 (3) (2014) 419–456. <https://doi.org/10.1007/s10208-013-9186-4>
- [39] G. Migliorati, F. Nobile, E. Von Schwerin, R. Tempone, Approximation of quantities of interest in stochastic PDEs by the random discrete  $L^2$  projection on polynomial spaces, *SIAM J. Sci. Comput.* 35 (3) (2013) A1440–A1460. <https://doi.org/10.1137/120897109>
- [40] G. Blatman, B. Sudret, Adaptive sparse polynomial chaos expansion based on least angle regression, *J. Comput. Phys.* 230 (6) (2011) 2345–2367. <https://doi.org/10.1016/j.jcp.2010.12.021>
- [41] J. Hampton, A. Doostan, Coherence motivated sampling and convergence analysis of least squares polynomial chaos regression, *Comput. Methods Appl. Mech. Eng.* 290 (2015) 73–97. <https://doi.org/10.1016/j.cma.2015.02.006>
- [42] A. Narayan, J.D. Jakeman, T. Zhou, A Christoffel function weighted least squares algorithm for collocation approximations, *Math. Comput.* 86 (306) (2017) 1913–1947. <https://doi.org/10.1090/mcom/3192>
- [43] A. Cohen, G. Migliorati, Optimal weighted least-squares methods, *SMAI J. Comput. Math.* 3 (2017) 181–203. <https://doi.org/10.5802/smai-jcm.24>
- [44] A. Cohen, M.A. Davenport, D. Leviatan, On the stability and accuracy of least squares approximations, *Found. Comput. Math.* 13 (5) (2013) 819–834. <https://doi.org/10.1007/s10208-013-9142-3>
- [45] M. Kohler, Nonlinear orthogonal series estimates for random design regression, *J. Stat. Plan. Inference* 115 (2) (2003) 491–520. [https://doi.org/10.1016/s0378-3758\(02\)00158-1](https://doi.org/10.1016/s0378-3758(02)00158-1)
- [46] L. Györfi, M. Kohler, A. Krzyżak, H. Walk, *A Distribution-Free Theory of Nonparametric Regression*, Springer New York, 2002. <https://doi.org/10.1007/b97848>
- [47] M. Croci, H.-A. Abdul-Lateef, I.C.J. Powell, An adaptive sampling algorithm for level-set approximation, [arXiv:2509.14896](https://arxiv.org/abs/2509.14896) (2025).
- [48] B. Adcock, A. Bao, J.D. Jakeman, A. Narayan, Compressed sensing with sparse corruptions: fault-tolerant sparse collocation approximations, *SIAM/ASA J. Uncertainty Quantif.* 6 (4) (2018) 1424–1453. <https://doi.org/10.1137/17M112590X>
- [49] Y. Shin, D. Xiu, Correcting data corruption errors for multivariate function approximation, *SIAM J. Sci. Comput.* 38 (4) (2016) A2492–A2511. <https://doi.org/10.1137/16M1059473>
- [50] M. Heene, A.P. Hinojosa, H.-J. Bungartz, D. Pflüger, A massively-parallel, fault-tolerant solver for high-dimensional PDEs, in: *Euro-Par 2016: Parallel Processing Workshops*, Springer International Publishing, 2017, p. 635–647. [https://doi.org/10.1007/978-3-319-58943-5\\_51](https://doi.org/10.1007/978-3-319-58943-5_51)
- [51] B. Harding, M. Hegland, J. Larson, J. Southern, Fault tolerant computation with the sparse grid combination technique, *SIAM J. Sci. Comput.* 37 (3) (2015) C331–C353. <https://doi.org/10.1137/140964448>
- [52] S. Pauli, P. Arbenz, C. Schwab, Intrinsic fault tolerance of multilevel Monte Carlo methods, *J. Parallel Distrib. Comput.* 84 (2015) 24–36. <https://doi.org/10.1016/j.jpdc.2015.07.005>
- [53] I. Babuška, F. Nobile, R. Tempone, A stochastic collocation method for elliptic partial differential equations with random input data, *SIAM J. Numer. Anal.* 45 (3) (2007) 1005–1034. <https://doi.org/10.1137/050645142>
- [54] C.W. Clenshaw, A.R. Curtis, A method for numerical integration on an automatic computer, *Numer. Math.* 2 (1) (1960) 197–205. <https://doi.org/10.1007/BF01386223>
- [55] V. Barthelmann, E. Novak, K. Ritter, High dimensional polynomial interpolation on sparse grids, *Adv. Comput. Math.* 12 (4) (2000) 273–288. <https://doi.org/10.1023/a:1018977404843>
- [56] M. Griebel, M. Schneider, C. Zenger, A combination technique for the solution of sparse grid problems, *Iterat. Methods Linear Algebra* (1992) 263–281. [https://www.google.co.uk/books/edition/Iterative\\_Methods\\_in\\_Linear\\_Algebra/H37CzwEACAAJ?hl=en](https://www.google.co.uk/books/edition/Iterative_Methods_in_Linear_Algebra/H37CzwEACAAJ?hl=en).
- [57] G.W. Wasilkowski, H. Woźniakowski, Explicit cost bounds of algorithms for multivariate tensor product problems, *J. Complex* 11 (1) (1995) 1–56. <https://doi.org/10.1006/jcom.1995.1001>
- [58] P.G. Constantine, M.S. Eldred, E.T. Phipps, Sparse pseudospectral approximation method, *Comput. Methods Appl. Mech. Eng.* 229–232 (2012) 1–12. <https://doi.org/10.1016/j.cma.2012.03.019>
- [59] L. Formaggia, A. Guadagnini, I. Imperiali, V. Lever, G. Porta, M. Riva, A. Scotti, L. Tamellini, Global sensitivity analysis through polynomial chaos expansion of a basin-scale geochemical compaction model, *Comput. Geosci.* 17 (1) (2012) 25–42. <https://doi.org/10.1007/s10596-012-9311-5>
- [60] F. Nobile, R. Tempone, C.G. Webster, An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data, *SIAM J. Numer. Anal.* 46 (5) (2008) 2411–2442. <https://doi.org/10.2307/25663069>
- [61] A.-L. Haji-Ali, F. Nobile, L. Tamellini, R. Tempone, Multi-index stochastic collocation convergence rates for random PDEs with parametric regularity, *Found. Comput. Math.* 16 (6) (2016) 1555–1605. <https://doi.org/10.1007/s10208-016-9327-7>
- [62] J. Zech, D. Düng, C. Schwab, Multilevel approximation of parametric and stochastic PDEs, *Math. Models Methods Appl. Sci.* 29 (09) (2019) 1753–1817. <https://doi.org/10.1142/s0218202519500349>
- [63] J. Beck, L. Tamellini, R. Tempone, IGA-Based multi-index stochastic collocation for random PDEs on arbitrary domains, *Comput. Methods Appl. Mech. Eng.* 351 (2019) 330–350. <https://doi.org/10.1016/j.cma.2019.03.042>
- [64] J.D. Jakeman, S. Friedman, M.S. Eldred, L. Tamellini, A.A. Gorodetsky, D. Allaire, Adaptive experimental design for multi-fidelity surrogate modeling of multi-disciplinary systems, *Int. J. Numer. Methods Eng.* 123 (12) (2022) 2760–2790. <https://doi.org/10.1002/nme.6958>
- [65] A.L. Teckentrup, P. Jantsch, C.G. Webster, M. Gunzburger, A multilevel stochastic collocation method for partial differential equations with random input data, *SIAM/ASA J. Uncertainty Quantif.* 3 (1) (2015) 1046–1074. <https://doi.org/10.1137/140969002>

- [66] A.A. Gorodetsky, G. Geraci, M.S. Eldred, J.D. Jakeman, A generalized approximate control variate framework for multifidelity uncertainty quantification, *J. Comput. Phys.* 408 (2020) 109257. <https://doi.org/10.1016/j.jcp.2020.109257>
- [67] Y. Xu, V. Keshavarzzadeh, R.M. Kirby, A. Narayan, A bandit-learning approach to multifidelity approximation, *SIAM J. Sci. Comput.* 44 (1) (2022) A150–A175. <https://doi.org/10.1137/21M1408312>
- [68] J.P. Boyd, Solving Transcendental Equations: the Chebyshev Polynomial Proxy and Other Numerical Rootfinders, Perturbation Series, and Oracles, Society for Industrial and Applied Mathematics, 2014. <https://doi.org/10.1137/1.9781611973525>
- [69] J.L. Aurentz, L.N. Trefethen, Chopping a Chebyshev series, *ACM Trans. Math. Softw.* 43 (4) (2017). <https://doi.org/10.1145/2998442>
- [70] P.-O. Persson, J. Peraire, Sub-cell shock capturing for discontinuous Galerkin methods, in: 44th AIAA Aerospace Sciences Meeting and Exhibit, American Institute of Aeronautics and Astronautics, 2006. <https://doi.org/10.2514/6.2006-112>
- [71] A. Huerta, E. Casoni, J. Peraire, A simple shock-capturing technique for high-order discontinuous Galerkin methods, *Int. J. Numer. Methods Fluids* 69 (10) (2011) 1614–1632. <https://doi.org/10.1002/flid.2654>
- [72] R. Killick, P. Fearnhead, I.A. Eckley, Optimal detection of changepoints with a linear computational cost, *J. Am. Stat. Assoc.* 107 (500) (2012) 1590–1598. <https://doi.org/10.1080/01621459.2012.737745>
- [73] A. Genz, Testing multidimensional integration routines, in: *Proc. of International Conference on Tools, Methods and Languages for Scientific and Engineering Computation*, Elsevier North-Holland, Inc., USA, 1984, p. 81–94.
- [74] H. Elman, D. Silvester, A. Wathen, *Finite Elements and Fast Iterative Solvers: with Applications in Incompressible Fluid Dynamics*, Oxford University Press, Oxford, UK, 2 edition, Oxford, UK, 2014. <https://doi.org/10.1093/acprof:oso/9780199678792.001.0001>
- [75] H.C. Elman, A. Ramage, D.J. Silvester, IFISS: a computational laboratory for investigating incompressible flow problems, *SIAM Rev.* 56 (2) (2014) 261–273. <https://doi.org/10.1137/120891393>
- [76] R.M. Smith, A.G. Hutton, The numerical treatment of advection: a performance comparison of the current methods, *Numer. Heat Transf.* 5 (4) (1982) 439–461. <https://doi.org/10.1080/10407788208913458>
- [77] K.W. Morton, *Numerical Solution of Convection-Diffusion Problems*, Chapman & Hall, London, UK, London, UK, 1996.
- [78] 2DWMH: 2D NASA wall-mounted hump separated flow validation case, 1996. accessed 2025-03-28, [https://turbmodels.larc.nasa.gov/nasahump\\_val.html](https://turbmodels.larc.nasa.gov/nasahump_val.html).
- [79] D. Greenblatt, K. Paschal, C. Yao, J. Harris, N. Schaeffler, A. Washburn, A separation control CFD validation test case. part 1: baseline & steady suction, in: 2nd AIAA Flow Control Conference, American Institute of Aeronautics and Astronautics, 2004. <https://doi.org/10.2514/6.2004-2220>
- [80] V. Tsiolakis, M. Giacomini, R. Sevilla, C. Othmer, A. Huerta, Parametric solutions of turbulent incompressible flows in OpenFOAM via the proper generalised decomposition, *J. Comput. Phys.* 449 (2022) 110802. <https://doi.org/10.1016/j.jcp.2021.110802>
- [81] B.M. Kent, *Efficient Approximation of Parametric Parabolic Partial Differential Equations*, PhD thesis, University of Manchester, 2023.
- [82] A. Iserles, *A First Course in the Numerical Analysis of Differential Equations*, Cambridge Texts in Applied Mathematics, Cambridge University Press, Cambridge, UK, 1 edition, Cambridge, UK, 1996. <https://doi.org/10.1017/CBO9780511995569>
- [83] P.M. Gresho, D.F. Griffiths, D.J. Silvester, Adaptive time-stepping for incompressible flow part I: scalar advection-diffusion, *SIAM J. Sci. Comput.* 30 (4) (2008) 2018–2054. <https://doi.org/10.1137/070688018>
- [84] F. Nobile, R. Tempone, Analysis and implementation issues for the numerical approximation of parabolic equations with random coefficients, *Int. J. Numer. Methods Eng.* 80 (6–7) (2009) 979–1006. <https://doi.org/10.1002/nme.2656>
- [85] J. Naughton, S. Viken, D. Greenblatt, Wall shear stress measurements on the NASA hump model for CFD validation, in: 24th AIAA Aerodynamic Measurement Technology and Ground Testing Conference, American Institute of Aeronautics and Astronautics, 2004. <https://doi.org/10.2514/6.2004-2607>
- [86] P. Spalart, S. Allmaras, A one-equation turbulence model for aerodynamic flows, in: 30th Aerospace Sciences Meeting and Exhibit, American Institute of Aeronautics and Astronautics, 1992. <https://doi.org/10.2514/6.1992-439>
- [87] The OpenFOAM foundation. <https://openfoam.org/>.
- [88] S.V. Patankar, D.B. Spalding, A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows, *Int. J. Heat Mass Transf.* 15 (10) (1972) 1787–1806. [https://doi.org/10.1016/0017-9310\(72\)90054-3](https://doi.org/10.1016/0017-9310(72)90054-3)
- [89] L. Seelinger, A. Reinarz, M.B. Lykkegaard, R. Akers, A.M.A. Alghamdi, D. Aristoff, W. Bangerth, J. Bénézech, M. Diez, K. Frey, J.D. Jakeman, J.S. Jørgensen, K.-T. Kim, B.M. Kent, M. Martinelli, M. Parno, R. Pellegrini, N. Petra, N.A.B. Riis, K. Rosenfeld, A. Serani, L. Tamellini, U. Villa, T.J. Dodwell, R. Scheichl, Democratizing uncertainty quantification, *J. Comput. Phys.* 521 (2025) 113542. <https://doi.org/10.1016/j.jcp.2024.113542>
- [90] J.D. Jakeman, L.A. Barba, J.R.R.A. Martins, T. O’Leary-Roseberry, Verification and validation for trustworthy scientific machine learning, 2025. [arXiv:2502.15496](https://arxiv.org/abs/2502.15496)