Enhancing adaptive sparse grid approximations and improving refinement strategies using adjoint-based a posteriori error estimates

Jakeman, J.D.^{a,1,*}, Wildey, T.^{a,1}

^a*Sandia National Laboratories, Albuquerque, NM 87185, United States*

Abstract

In this paper we present an algorithm for adaptive sparse grid approximations of quantities of interest computed from discretized partial differential equations. We use adjoint-based a posteriori error estimates of the physical discretization error and the interpolation error in the sparse grid to enhance the sparse grid approximation and to drive adaptivity of the sparse grid. Utilizing these error estimates provides significantly more accurate functional values for random samples of the sparse grid approximation. We also demonstrate that alternative refinement strategies based upon a posteriori error estimates can lead to further increases in accuracy in the approximation over traditional hierarchical surplus based strategies. Throughout this paper we also provide and test a framework for balancing the physical discretization error with the stochastic interpolation error of the enhanced sparse grid approximation.

Keywords: Uncertainty quantification, a posteriori error estimation, sparse grids, stochastic collocation, adaptivity

Partial Differential Equations (PDE) are used to simulate a wide range of phenomenon and are often used to inform design decisions and to estimate risk in systems with large human and/or financial impact but with limited capacity for experimentation. Given the importance of these appli-

*Corresponding author

Email address: jdjakem@sandia.gov (Jakeman, J.D.)

¹Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

cations the ability to accurately quantify uncertainty in model predictions is essential.

Most uncertainty quantification (UQ) studies focus on estimating parametric uncertainty. In such analyses, the uncertainty in the input data, such as model coefficients, forcing terms etc, is usually represented through a finite number of random variables with a known probability distribution. The goal of the study is then to compute the effect of the varying input data on the system response, and in many cases, to calculate the statistics of the response.

The accuracy to which uncertainty can be quantified is limited by the computational resources available to resolve the governing equations. Many models require vast amounts of computational effort and thus the number of model evaluations that can be used to interrogate the uncertainty in the system behavior is limited. Consequently a significant portion of methods developed for uncertainty quantification (UQ) in recent years have focused on constructing surrogates of expensive simulation models using only a limited number of model evaluations.

The most widely adopted approximation methods are based on generalized polynomial chaos (PC) expansions [14, 26], sparse grid interpolation [19, 20] and Gaussian Process (GP) models [24]. The performance of these methods is problem dependent and in practice it is difficult to estimate the accuracy of the approximation constructed. Cross-validation is one means of estimating the accuracy of the approximation, however the accuracy of the cross-validation prediction of the error is limited. Moreover, cross validation is not readily applied for approximation methods which require structured model samples, such as sparse grid interpolation and many forms of pseudo-spectral projection.

In this paper we utilize sparse grid interpolation to approximate model responses. Sparse grids can be built using local or global basis functions and have well established and effective adaptivity procedures which can be leveraged in conjunction with good error estimates to concentrate computational effort to resolving important dimensions and/or regions of the random parameter space. Unlike regression based PCE or Gaussian Process models, sparse grids can be used regardless of the computational budget of the UQ analysis. For example sparse grids can be used to approximate a model response using tens to millions of model runs, whereas the aforementioned alternatives have upper limits in the low thousands imposed by the need to solve large linear systems.

Throughout this paper, we will use $J(\boldsymbol{\xi})$ to denote the exact response (quantity of interest) from a partial differential equation that depends on

the unknown variable $\boldsymbol{\xi}$. When solving PDEs using techniques such as the finite element method the physical discretization error will be non-zero. We use $J_h(\boldsymbol{\xi})$ to denote the response from the discretized model. As previously mentioned, solving the discretized model is often computationally expensive and therefore we need to consider a surrogate approximation of $J_h(\boldsymbol{\xi})$, which we denote $J_{h,n}(\boldsymbol{\xi})$. Given these approximations, the error in the response can be decomposed into two components

$$\|J(\boldsymbol{\xi}) - J_{h,n}(\boldsymbol{\xi})\| \leq \underbrace{\|J(\boldsymbol{\xi}) - J_h(\boldsymbol{\xi})\|}_I + \underbrace{\|J_h(\boldsymbol{\xi}) - J_{h,n}(\boldsymbol{\xi})\|}_{II} \quad (0.1)$$

where: I is the finite element *physical discretization* error; and II is the *stochastic approximation* error introduced by approximating the quantity of interest by a sparse grid interpolant.

Recently, a posteriori error estimation has arisen as a promising approach to estimate the error in approximate input-output relationships. Adjoint-based a posteriori error estimation was originally developed to estimate error in numerical approximations of deterministic Partial Differential Equations (PDE) [2, 10, 15, 23], but recent modifications allow similar ideas to be used to estimate error in approximations of quantities of interest from PDEs with uncertain parameters. This relatively new approach, introduced in [7] and further analyzed in [6, 8], is based on goal-oriented adjoint-based error estimates and is used to predict error in samples of a response surface approximation of a specific quantity of interest. Similar to standard adjoint-based error estimation procedures, this new approach includes the physical discretization error if the adjoint problem is approximated in a higher-order discretization space. However, the error estimate from this new approach also contains an approximation of the error in the stochastic discretization due to the evaluation of the response surface model rather than the PDE. In [6, 8], it was shown that, for spectral and pseudo-spectral Galerkin approximations, this estimate of the stochastic interpolation error is higher-order even if a low order approximation of the adjoint is used for the stochastic approximation.

In general, it is inefficient to reduce the stochastic error to a level below the error introduced by the deterministic discretization. Much of the existing literature focuses on minimizing the stochastic approximation error, however only a few attempts have been made to discuss or account for the combined effect of deterministic and stochastic approximation error. Error bounds for the stochastic approximation error for isotropic sparse grid approximations of elliptic PDEs using Clenshaw-Curtis or Gaussian abscissa are given in

[21]. In this paper, we use adjoint-based error estimates to ensure that the error in the stochastic approximation is never significantly less than the physical discretization error.

Our goal in this paper is to utilize adjoint-based a posteriori error estimates to efficiently compute pointwise approximations of specific quantities of interest, usually computed from PDE solutions, using adaptive sparse grid approximations. Specifically, we aim to

- Provide theoretical bounds on the error in a posteriori enhanced Clenshaw-Curtis sparse grids.
- Numerically demonstrate the enhancement results in [6] extend to adaptive sparse grid approximations.
- Present new refinement strategies for sparse grids based on a posteriori error estimates.
- Present a strategy for reducing the cost of computing a posteriori error estimates.
- Provide a criterion to stop sparse grid refinement when the stochastic approximation error of the sparse grid is approximately equal to the physical discretization error.

The remainder of this paper is organized as follows. Section 1 introduces the general model problem we are interested in. Sparse grid approximation is reviewed in Section 2 and we recall the standard adjoint-based posteriori error analysis for deterministic PDEs in Section 3. In Section 4 we formulate an a posteriori error estimate for samples of a sparse grid surrogate and derive theoretical bounds on the error in the a posteriori error estimate. Section 5, presents new adaptive strategies for sparse grid refinement that leverage a posteriori error estimates and Section 6 introduces the sparse grid approximation of the error estimate and our stopping criteria based on an estimate the physical discretization error. Numerical results are presented in Section 7 and our conclusions are presented in Section 8.

1. General nonlinear problem and notation

We consider the following system,

$$\frac{\partial \mathbf{z}}{\partial t} + \mathbf{A}(\boldsymbol{\xi}; \mathbf{z}) = \mathbf{0}. \quad (1.1)$$

In this paper, we consider systems of ordinary differential equations (ODEs) as well as systems of partial differential equations (PDEs). We focus most of our attention on the numerical approximation of PDEs since the corresponding analysis for the ODE case is a straightforward simplification. In the PDE case, (1.1) is defined on $\Omega \times (0, T]$ where $\Omega \subset \mathbb{R}^s$, $s = 1, 2, 3$, is a polygonal (polyhedral) and bounded domain with boundary $\partial\Omega$. The random parameter $\boldsymbol{\xi}$ takes values in $\Gamma_{\boldsymbol{\xi}} \subset \mathbb{R}^d$ and reflects uncertainty in model and source parameters. The solution operator's dependency on $\boldsymbol{\xi}$ implies that $\mathbf{z} := \mathbf{z}(\mathbf{x}, t, \boldsymbol{\xi})$ is also uncertain and may be modeled as a random process for which we will construct a sparse grid approximation.

Before introducing assumptions on the properties of $\mathbf{A}$ and $\mathbf{z}$ let us first introduce some notation. Let $(\cdot, \cdot)_{\mathcal{D}}$ denote the inner product of $\mathbf{L}^2(\mathcal{D})$ and if the domain of integration is clear from the context, we suppress the index $\mathcal{D}$. Define $\mathbf{V}$ be a Sobolev space where for any non-negative integer m we recall

$$H^m(\Omega) = \{v \in L^2(\Omega); \partial^k v \in L^2(\Omega) \forall |k| \leq m\},$$

equipped with the following seminorm and norm

$$|v|_{H^m(\Omega)} = \left[\sum_{|k|=m} \int_{\Omega} |\partial^k v|^2 dx \right]^{1/2}, \quad \|v\|_{H^m(\Omega)} = \left[\sum_{0 \leq |k| \leq m} |v|_{H^k(\Omega)}^2 \right]^{1/2}.$$

We use $\mathbf{H}^m(\Omega)$ to denote the obvious generalization to vector valued functions.

In this paper we assume that $\mathbf{A}$ is convex and has smooth second derivatives. Specific examples of $\mathbf{A}$ and $\mathbf{z}$ will be given in subsequent sections. We assume that sufficient initial and boundary conditions are provided so that (1.1) is well-posed in the sense that there exists a solution for a. e. $\boldsymbol{\xi} \in \Gamma_{\boldsymbol{\xi}}$. Specifically we assume that (1.1) has an equivalent variational formulation seeking $\mathbf{z} \in \mathbf{V}$, for any $\boldsymbol{\xi} \in \Gamma_{\boldsymbol{\xi}}$, such that

$$\int_0^T \left[\left(\frac{\partial \mathbf{z}}{\partial t}, \mathbf{w} \right) + \mathbf{a}(\boldsymbol{\xi}; \mathbf{z}, \mathbf{w}) \right] dt = \mathbf{0}, \quad \forall \mathbf{w} \in \mathbf{V}, 0 < t \leq T, \quad (1.2)$$

where $\mathbf{a}(\boldsymbol{\xi}; \cdot, \cdot)$ is the weak form of $\mathbf{A}$.

Now let $\mathcal{T}_h$ be a conforming partition of Ω , composed of N_T closed convex volumes of maximum diameter h . An element of the partition $\mathcal{T}_h$ will be denoted by T_i where h_i stands for the diameter of T_i for $i = 1, 2, \dots, N_T$. We assume that the mesh is regular in the sense of Ciarlet [9] and take $\mathcal{T}_h$ to be a conforming finite element mesh consisting of simplices or parallelopipeds.

Denoting the space of continuous, piecewise functions of degree q (simplices) or total degree q (parallelopipeds) over the spatial domain by

$$\mathbf{V}_h^{(q)} = \{ \mathbf{v} \in C(\Omega) \cap \mathbf{H}^1(\Omega) : \forall E \in \mathcal{T}_h, \mathbf{v}|_E \in \mathbb{P}^q(E) \}.$$

we then use a semi-discrete variational formulation to seek $\mathbf{z}_h \in \mathbf{V}_h^{(q)}$, for any $\boldsymbol{\xi} \in \Gamma_{\boldsymbol{\xi}}$, such that

$$\int_0^T \left[\left(\frac{\partial \mathbf{z}_h}{\partial t}, \mathbf{w} \right) + \mathbf{a}(\boldsymbol{\xi}; \mathbf{z}_h, \mathbf{w}) \right] dt = \mathbf{0}, \quad \forall \mathbf{w} \in \mathbf{V}_h^{(q)}, 0 < t \leq T. \quad (1.3)$$

A fully discrete scheme for any $\boldsymbol{\xi} \in \Gamma_{\boldsymbol{\xi}}$ can be obtained by letting $I_n = (t_{n-1}, t_n)$ and time steps $k_n = t_n - t_{n-1}$ denote the discretization of $[0, T]$ as $0 = t_0 < t_1 < \dots < t_N = T$. In this paper, we do not focus on any particular time integration scheme. We do assume that the fully discrete approximation is given as a polynomial in time, e.g.,

$$\mathbf{z}_h \in \mathbf{V}_h^{(q)} \times \mathcal{W}^{(r)},$$

where

$$\mathcal{W}^{(r)} = \{ w \in C([0, T]) : w \in \mathbb{P}^r(I_n), \forall I_n \},$$

denotes the space of continuous piecewise polynomial functions of degree r . To make the notation less cumbersome, we generally use $\mathbf{z}_h$ to denote the fully discrete solution unless otherwise noted. The choice of continuous polynomials in time is merely for convenience. Discontinuous polynomials may also be used and require a straightforward modification to the a posteriori error analysis involving jump terms at each time node [10, 11].

In cases where a unique and sufficiently regular solution to (1.2) exists, one can obtain the following error bound (see e.g. [4, 17, 11, 10] for specific examples) for any $\boldsymbol{\xi} \in \Gamma_{\boldsymbol{\xi}}$,

$$\|\mathbf{z} - \mathbf{z}_h\|_{L^\infty((0, T); \mathbf{H}^1(\Omega))} \leq C(\mathbf{z})(h^s + k^\beta), \quad (1.4)$$

where $C(\mathbf{z})$ depends but does not depend on $k = \max_n k_n$ or h . The parameters, s and β , are determined by the regularity of the solution and the order of accuracy of the spatial and temporal discretizations. We emphasize that obtaining error bounds of this sort is nontrivial and may not be possible in many cases. Throughout the remainder of this paper we assume that (1.2) does have a unique and sufficiently regular solution to obtain (1.4).

2. Sparse grid approximations

Sparse grid stochastic collocation has been shown to provide efficient and accurate approximation of stochastic quantities [19, 20, 21]. Here we adopt this technique to construct approximations of functionals of the discretized solution, $J_h(\boldsymbol{\xi}) := J(\mathbf{z}_h(\boldsymbol{\xi}))$. We typically do not know the closed form of $J_h(\boldsymbol{\xi})$ and only require that we can evaluate $J_h(\boldsymbol{\xi})$ at arbitrary points in $\Gamma_{\boldsymbol{\xi}}$ by evaluating the discretized model. For simplicity we will restrict attention to consider stochastic collocation problems characterized by variables $\boldsymbol{\xi}$ with finite support normalized to fit in the domain $\Gamma_{\boldsymbol{\xi}} = [0, 1]^d$. However the technique proposed here can be applied to semi or unbounded random variables using the methodology outlined in [18].

Sparse grids [5] approximate $J_h(\boldsymbol{\xi})$ via a weighted linear combination of basis functions

$$I_n[J_h(\boldsymbol{\xi})] := J_{h,n}(\boldsymbol{\xi}) = \sum_{k=1}^n v_k \Psi_k(\boldsymbol{\xi}) \quad (2.1)$$

The approximation is constructed on a set of anisotropic grids $\Xi_{\mathbf{l}}$ on the domain $\Gamma_{\boldsymbol{\xi}}$ where $\mathbf{l} = (l_1, \dots, l_d) \in \mathbb{N}^d$ is a multi-index denoting the level of refinement of the grid in each dimension. These rectangular grids are Cartesian product of nested one-dimensional grid points $\Xi_l = \{\xi_{l,i} : i < 0 \leq i \leq m_l\}$

$$\Xi_{\mathbf{l}} = \Xi_{l_1} \times \dots \times \Xi_{l_d}$$

Typically when approximating $J_h(\boldsymbol{\xi})$ with a smooth dependence on $\boldsymbol{\xi}$, Ξ_l are chosen to be the nested Gaussian quadrature rules associated with the distribution of ξ_i . For example the Gauss-Patterson rule is used for uniform variables and Genz-Keister rule for Gaussian variables. For functions of lower regularity Ξ_l are typically chosen to be equidistantly spaced. The number of points m_l of a one-dimensional grid of a given level is dependent on the growth rate of the quadrature rule chosen.

The multivariate basis functions Ψ_k are a tensor product of one dimensional basis functions. Adopting the multi-index notation use above we have

$$\Psi_{\mathbf{l}, \mathbf{i}}(\boldsymbol{\xi}) = \prod_{n=1}^d \psi_{l_n, i_n}(\xi_n) \quad (2.2)$$

where $\mathbf{i}$ determines the location of a given grid point. There is a one-to-one relationship between Ψ_k in (2.1) and $\Psi_{\mathbf{l}, \mathbf{i}}$ and each $\Psi_{\mathbf{l}, \mathbf{i}}$ is uniquely associated with a grid point $\boldsymbol{\xi}_{\mathbf{l}, \mathbf{i}} = (\xi_{l_1, i_1}, \dots, \xi_{l_d, i_d}) \in \Xi_{\mathbf{l}}$. Many different one-dimensional basis functions $\psi_{l_n, i_n}(\xi_n)$ can be used. If $J_h(\boldsymbol{\xi})$ has a smooth

dependence on $\boldsymbol{\xi}$ then the best choice is the one-dimensional Lagrange polynomials. If local approximation is required one can use a multi-dimensional piecewise polynomial basis [5].

The multi-dimensional basis (2.2) span the discrete spaces $V_1 \subset L^2(\Gamma_\xi)$

$$V_1 = \text{span} \{ \Psi_{1,\mathbf{i}} : \mathbf{i} \in \mathcal{K}_1 \} \quad \mathcal{K}_1 = \{ \mathbf{i} : i_k = 0, \dots, m_{l_k}, k = 1, \dots, d \}$$

These discrete spaces can be further decomposed into hierarchical difference spaces

$$W_1 = V_1 \setminus V_1 \bigoplus_{n=0}^d V_{1-\mathbf{e}_n}$$

The subspaces W_1 consists of all basis functions $\Psi_{1,\mathbf{i}} \in V_1$ which are not included in any of the spaces $V_{\mathbf{k}}$ smaller than V_1 , i.e. with $\mathbf{k} < \mathbf{1}$. These hierarchical difference spaces can be used to decompose the input space such that

$$V_1 = \bigoplus_{\mathbf{k} \leq \mathbf{1}} W_1 \quad \text{and} \quad L^2(\Gamma_\xi) = \bigoplus_{k_1=0}^{\infty} \cdots \bigoplus_{k_d=0}^{\infty} W_{\mathbf{k}} = \bigoplus_{\mathbf{k} \in \mathbb{R}^d} W_{\mathbf{k}}$$

For numerical purposes we must truncate the number of difference spaces used to construct V . Traditional isotropic sparse grids can be obtained by all hierarchical subspaces W_1 with an index set that satisfy

$$\mathcal{L} = \{ \mathbf{l} : |\mathbf{l}|_1 \leq l \} \quad (2.3)$$

Given a truncation, such as the a-priori one above or one which has been determined adaptively, $J_h(\boldsymbol{\xi})$ can be approximated by

$$J_{h,n}(\boldsymbol{\xi}) = \sum_{\mathbf{l} \in \mathcal{L}} J_{\mathbf{l}}, \quad J_{\mathbf{l}} = \sum_{\mathbf{i} \in \mathcal{I}_{\mathbf{l}}} v_{\mathbf{l},\mathbf{i}} \Psi_{\mathbf{l},\mathbf{i}}(\boldsymbol{\xi}) \quad (2.4)$$

where $\mathcal{I}_{\mathbf{l}} = \{ \mathbf{i} : \Psi_{\mathbf{l},\mathbf{i}} \in W_{\mathbf{l}} \}$.

Here we note that the $v_{\mathbf{l},\mathbf{i}}$ are the coefficient values of the hierarchical product basis, also known as the hierarchical surplus. The surpluses are simply the difference between the function value and the sparse grid approximation at a point, not already in the sparse grid. That is

$$v_{\mathbf{l},\mathbf{i}} = J(\boldsymbol{\xi}_{\mathbf{l},\mathbf{i}}) - J_n(\boldsymbol{\xi}_{\mathbf{l},\mathbf{i}}), \quad \mathcal{L} \cup \mathbf{1} = \emptyset$$

3. Adjoint-based a posteriori error estimation

The goal of a simulation is often to accurately estimate a relatively small number of quantities of interest $J(\boldsymbol{\xi})$. Adjoint-based error analysis relates the error ($\mathbf{e} = \mathbf{z} - \mathbf{z}_h$) in a quantity of interest to a computable weighted residual. We are usually interested in estimating the error in a numerical approximation computed using a discretization of a variational formulation, but it is often intuitive to start the discussion with strong form adjoint operators. We assume that we are interested in estimating the error in a linear functional of the solution,

$$J(\boldsymbol{\xi}) = (\boldsymbol{\psi}_T(x), \mathbf{z}(x, T, \boldsymbol{\xi})) + \int_0^T (\boldsymbol{\psi}, \mathbf{z}(x, t, \boldsymbol{\xi})) dt, \quad (3.1)$$

where $\boldsymbol{\psi}_T(x)$ is used to compute a linear functional at $t = T$, and $\boldsymbol{\psi} = \boldsymbol{\psi}(x, t)$ is used to compute time-averages. Typically, either $\boldsymbol{\psi}_T$ or $\boldsymbol{\psi}$ are zero, but this is not required. Nonlinear functionals and linear functionals at $t \neq T$ can also be used, but this is omitted here for the sake of simplicity.

For a given $\boldsymbol{\xi} \in \Gamma_{\boldsymbol{\xi}}$, the linear adjoint operator in strong form is defined via the duality relation

$$\int_0^T \left[\left(\frac{\partial \mathbf{z}}{\partial t}, \mathbf{v} \right) + (\mathcal{L}(\boldsymbol{\xi})\mathbf{v}, \mathbf{w}) \right] dt = \int_0^T \left[\left(-\frac{\partial \mathbf{v}}{\partial t}, \mathbf{z} \right) + (\mathbf{v}, \mathcal{L}^*(\boldsymbol{\xi})\mathbf{w}) \right] dt \quad (3.2)$$

where $\mathcal{L}(\boldsymbol{\xi})$ is a given linear operator. For a general nonlinear PDE one approach to define the linear operator $\mathcal{L}(\boldsymbol{\xi})$ is to assume $\mathbf{A}(\boldsymbol{\xi}; \cdot)$ is convex and use the Integral Mean Value Theorem yielding

$$\mathbf{A}'(\boldsymbol{\xi}, \bar{\mathbf{z}}; \mathbf{e}) = \mathbf{A}(\boldsymbol{\xi}; \mathbf{z}) - \mathbf{A}(\boldsymbol{\xi}; \mathbf{z}_h)$$

where $\bar{\mathbf{z}}$ lies on the line connecting $\mathbf{z}$ and $\mathbf{z}_h$, and $\mathbf{e} = \mathbf{z} - \mathbf{z}_h$. In practice, $\bar{\mathbf{z}}$ is unknown so we linearize around $\mathbf{z}_h$ giving,

$$\mathcal{L}(\boldsymbol{\xi})\mathbf{e} = \mathbf{A}'(\boldsymbol{\xi}, \mathbf{z}_h; \mathbf{e}) = \mathbf{A}(\boldsymbol{\xi}; \mathbf{z}) - \mathbf{A}(\boldsymbol{\xi}; \mathbf{z}_h) + \mathbf{R}(\boldsymbol{\xi}; \mathbf{e}, \mathbf{z}_h, \bar{\mathbf{z}}),$$

where $\mathbf{R}(\boldsymbol{\xi}; \mathbf{e}, \mathbf{z}_h, \bar{\mathbf{z}})$ represents the remainder. Since $\mathbf{z}_h - \bar{\mathbf{z}} \approx \mathbf{e}$, it is common to assume that the remainder is a higher order perturbation term and can be neglected [2, 3, 11]. Notice that the operator $\mathcal{L}(\boldsymbol{\xi})$ is often the same linear operator used in computing the step in Newton's method. This fact is often exploited to ease construction of the discrete adjoint operator.

The error in a linear functional can be represented using the definition of the adjoint:

$$J(\boldsymbol{\xi}) - J_h(\boldsymbol{\xi}) = (\phi(x, 0, \boldsymbol{\xi}), \mathbf{z}(x, 0, \boldsymbol{\xi}) - \mathbf{z}_h(x, 0, \boldsymbol{\xi})) - \int_0^T \left(\left(\frac{\partial \mathbf{z}_h}{\partial t}, \phi \right) + \mathbf{a}(\boldsymbol{\xi}; \mathbf{z}_h, \phi) \right) dt + \int_0^T \mathbf{R}(\boldsymbol{\xi}; \mathbf{e}, \mathbf{z}_h, \bar{\mathbf{z}}, \phi) dt, \quad (3.3)$$

where $\phi := \phi(x, t, \boldsymbol{\xi})$ is defined by the adjoint problem

$$-\frac{\partial \phi}{\partial t} + \mathcal{L}^*(\boldsymbol{\xi})\phi = \psi(x, t), \quad (3.4)$$

$$\phi(x, T) = \psi_T(x). \quad (3.5)$$

If the adjoint solution, ϕ , is given, then the error representation in Eq. (3.3) is easily evaluated if we neglect the higher order remainder term, see [2, 3, 11, 12] for a complete discussion of this remainder term. However, the solution to the adjoint problem Eq. (3.4) is usually not given explicitly and we must approximate the solution using an appropriate discretization. In this paper, we approximate the adjoint solution using a finite element method. Since we are also interested in estimating the physical discretization error, we use a higher-order approximation in the spatial domain. This gives a fully computable error representation:

$$\varepsilon(\mathbf{z}_h, \phi_h) = (\phi_h(x, 0, \boldsymbol{\xi}), \mathbf{z}(x, 0, \boldsymbol{\xi}) - \mathbf{z}_h(x, 0, \boldsymbol{\xi})) - \int_0^T \left(\left(\frac{\partial \mathbf{z}_h}{\partial t}, \phi_h \right) + \mathbf{a}(\boldsymbol{\xi}; \mathbf{z}_h, \phi_h) \right) dt. \quad (3.6)$$

In general, the definition of the adjoint problem also requires appropriate boundary conditions and the error representation (3.3) also includes boundary terms to account for errors made in approximating the boundary conditions on the forward problem. We omit these terms in this paper for the sake of simplicity and refer the interested reader to the standard references, e.g. [2, 10, 15, 23], for a complete discussion of these additional terms.

4. Adjoint-based error estimates for samples of a sparse grid surrogate

In many practical situations, we are interested in computing probabilistic quantities, such as the probability of a particular event, from the surrogate

approximation. Computing probabilistic quantities usually requires sampling the surrogate according to the distribution of the random parameters over $\Gamma_{\boldsymbol{\xi}}$. In order to have confidence in our estimates of these probabilistic quantities, we must consider the error in each sample of the surrogate.

In this paper, we employ the technique introduced in [7] and further analyzed in [6, 8] to estimate the error in a sample of the sparse grid surrogate of the quantity of interest.

Given sparse grid approximations of the forward $\mathbf{z}_{h,n}$ and adjoint $\phi_{h,n}$ solutions we can compute the following approximate error estimate

$$J(\boldsymbol{\xi}) - J_{h,n}(\boldsymbol{\xi}) \approx \varepsilon(\mathbf{z}_{h,n}, \phi_{h,n}), \quad (4.1)$$

where

$$\begin{aligned} \varepsilon(\mathbf{z}_{h,n}, \phi_{h,n}) &= (\phi_{h,n}(x, 0, \boldsymbol{\xi}), \mathbf{z}(x, 0, \boldsymbol{\xi}) - \mathbf{z}_{h,n}(x, 0, \boldsymbol{\xi})) \\ &- \int_0^T \left(\left(\frac{\partial \mathbf{z}_{h,n}(x, t, \boldsymbol{\xi})}{\partial t}, \phi_{h,n}(x, t, \boldsymbol{\xi}) \right) + \mathbf{a}(\boldsymbol{\xi}; \mathbf{z}_{h,n}(x, t, \boldsymbol{\xi}), \phi_{h,n}(x, t, \boldsymbol{\xi})) \right) dt. \end{aligned}$$

The error in $\varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})$ is dependent on the errors in the approximation $\mathbf{z}_h$ and ϕ_h which are in turn dependent on the regularity of $\mathbf{z}$ and ϕ with respect to the random variables $\boldsymbol{\xi}$. For the analysis presented in this paper we will make the following regularity assumption.

Assumption 4.1 (Regularity Assumption). Let Γ_{ξ}^1 be the one dimension range of a random variable ξ and

$$\Gamma_{\boldsymbol{\xi}} = \prod_{i=1}^d \Gamma_{\xi_i}^1, \quad \Gamma_{\boldsymbol{\xi}}^{(-n)} = \prod_{\substack{i=1 \\ i \neq n}}^d \Gamma_{\xi_i}^1$$

For each $\xi_n \in \Gamma_n^1$ there exists $\tau_n > 0$ such that the function $\mathbf{z}_h(\xi_n, \boldsymbol{\xi}^{(-n)}, \mathbf{x}) : \Gamma_{\xi_n}^1 \rightarrow C^0(\Gamma_{\boldsymbol{\xi}}^{(-n)}; L^\infty(0, T) \times \mathbf{H}^1(\Omega))$ as a function of ξ_n , admits an analytical extension $\mathbf{z}(y, \boldsymbol{\xi}^{(-n)}, \mathbf{x})$, $y \in \mathbb{C}$ in the region of the complex plain

$$\Sigma(\Gamma_{\xi_n}^1; \tau_n) \equiv \{y \in \mathbb{C}, \text{dist}(y, \Gamma_{\xi_n}^1) \leq \tau_n\},$$

where $\boldsymbol{\xi}^{(-n)} \in \Gamma_{\boldsymbol{\xi}}^{(-n)}$. Moreover $\forall y \in \Sigma(\Gamma_{\xi_n}^1; \tau_n)$

$$\|\mathbf{z}_h(y)\|_{C^0(\Gamma_{\boldsymbol{\xi}}^{(-n)}; L^\infty(0, T) \times \mathbf{H}^1(\Omega))} \leq \lambda$$

for a constant λ independent of n .

This assumption is valid for a number of problems. Indeed the the linear elliptic diffusion equation analyzed in Section 4 satisfies assumption 4.1 [1]. Using this assumption we can obtain the following result.

Lemma 4.1 ([22]). For functions $\mathbf{z} \in C^0(\Gamma_\xi; L^\infty(0, T) \times \mathbf{H}^1(\Omega))$ satisfying Assumption 4.1 the isotropic Smolyak formula (2.8) based on Clenshaw-Curtis abscissas, and with n points, satisfies:

$$\|\mathbf{z}_h - \mathbf{z}_{h,n}\|_{L^\infty(\Gamma_\xi; L^\infty(0, T) \times \mathbf{H}^1(\Omega))} \leq C_1(\sigma_1) n^{-\mu_1}$$

$$\mu_1 := \frac{\sigma_1}{1 + \log 2d}$$

where the constant $C_1(\sigma)$ depends on the size of the analyticity region σ_1 but not on the number of points in the sparse grid.

Combining Lemma 4.1 and taking the maximum of (1.4) over Γ_ξ yields the following result.

Lemma 4.2 ([22]). For functions $\mathbf{z} \in C^0(\Gamma_\xi; L^\infty(0, T) \times \mathbf{H}^1(\Omega))$ satisfying Assumption 4.1, isotropic Smolyak formula (2.8) based on Clenshaw-Curtis abscissas satisfies:

$$\|\mathbf{z} - \mathbf{z}_{h,n}\|_{L^\infty(\Gamma_\xi; L^\infty(0, T) \times \mathbf{H}^1(\Omega))} \leq C_1(\sigma_1) n^{-\mu_1} + C_1(\mathbf{z}) (h^{s_1} + k^{\beta_1})$$

where the constant $C_1(\sigma_1)$ depends on the size of the analyticity region σ_1 but not on the number of points in the sparse grid and $C_1(\mathbf{z})$ depends on the solution $\mathbf{z}$ but not h or k .

We can also apply Lemmas 4.1 and 4.2 to the adjoint solution, ϕ , and the numerical approximation of the adjoint solution, $\phi_{h,n}$, giving

$$\|\phi - \phi_{h,n}\|_{L^\infty(\Gamma_\xi; L^\infty(0, T) \times \mathbf{H}^1(\Omega))} \leq C_2(\sigma_2) n^{-\mu_2} + C_2(\phi) (h^{s_2} + k^{\beta_2}). \quad (4.2)$$

Note that we account for the fact that the constants and convergence parameters may be different for the adjoint solution.

Next, we prove that the error estimate, $\varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})$, converges at a faster rate than $J_{h,n}(\xi)$. First, we note that we can approximate

$$J(\xi) - J_{h,n}(\xi) \approx \varepsilon(\mathbf{z}_{h,n}, \phi).$$

This approximation is exact for linear problems. For nonlinear problems, the error in this approximation is a high-order linearization term and is usually neglected. Thus, we prove that $\varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})$ is a high-order approximation of $\varepsilon(\mathbf{z}_{h,n}, \phi)$.

Theorem 4.3. For functions $\mathbf{z}, \phi \in C^0(\Gamma_\xi; L^\infty(0, T) \times \mathbf{H}^1(\Omega))$ satisfying Assumption 4.1 and using isotropic Smolyak formula (2.8) based on Clenshaw-Curtis abscissas, the error estimate satisfies:

$$\begin{aligned} \|\varepsilon(\mathbf{z}_{h,n}, \phi) - \varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})\|_{L^\infty(\Gamma_\xi)} &\leq \left(C_1(\sigma_1) n^{-\mu_1} + C_1(\mathbf{z}) \left(h^{s_1} + k^{\beta_1} \right) \right) \\ &\quad \times \left(C_2(\sigma_2) n^{-\mu_2} + C_2(\phi) \left(h^{s_2} + k^{\beta_2} \right) \right) \end{aligned}$$

where the constants $C(\sigma)$ and $C_2(\sigma_2)$ depend on the size of the analyticity regions σ_1 and σ_2 respectively, but not on the number of points in the sparse grid. In addition, $C_1(\mathbf{z})$ and $C_2(\phi)$ depend on the solutions $\mathbf{z}$ and ϕ respectively, but not h or k .

Proof. We use the fact that $\varepsilon(\cdot, \cdot)$ is linear in the second argument to write,

$$|\varepsilon(\mathbf{z}_{h,n}, \phi) - \varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})| = |\varepsilon(\mathbf{z}_{h,n}, \phi - \phi_{h,n})|.$$

It is straightforward (see e.g. [8, 23, 2]) to show that

$$\begin{aligned} |\varepsilon(\mathbf{z}_{h,n}, \phi - \phi_{h,n})| &\leq \|\mathbf{z} - \mathbf{z}_{h,n}\|_{L^\infty(\Gamma_\xi; L^\infty(0, T) \times \mathbf{H}^1(\Omega))} \\ &\quad \times \|\phi - \phi_{h,n}\|_{L^\infty(\Gamma_\xi; L^\infty(0, T) \times \mathbf{H}^1(\Omega))}. \end{aligned}$$

A direct application of Lemma 4.2 and (4.2) proves the result. $\square$

We are often interested in using $\varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})$ to improve the approximation of the quantity of interest. To this end, we define the enhanced quantity of interest to be

$$J_n^\varepsilon(\xi) := J_{h,n}(\xi) + \varepsilon(\mathbf{z}_{h,n}, \phi_{h,n}).$$

Theorem 4.3 implies that $J_n^\varepsilon(\xi)$ is higher-order approximation to $J(\xi)$ than $J_{h,n}(\xi)$.

5. Sparse grid adaptivity

Although a vast improvement of full tensor grid approximations the number of sparse grid points in the isotropic sparse grids, presented in Section 2 and analyzed in Section 4, still grows quickly with dimension. To enhance the convergence of sparse grids two types of adaptivity have been developed. The first type of adaptivity refines the grid dimension by dimension [13], greedily choosing points in dimensions that are deemed by the algorithm to be more important. The second type of adaptivity refines the sparse grid locally in regions identified as important [16]. When using localized basis functions adaptivity can be performed combining the strengths of both dimension and regional adaptivity [19].

5.1. Dimension adaptivity

The dimension adaptive algorithm begins with a low-level isotropic sparse grid approximation with a set of subspaces $\mathcal{L}$ and active subspaces $\mathcal{A}$. Often $\mathcal{L} = W_{\mathbf{0}}$ and $\mathcal{A} = \{W_{\mathbf{e}_k}, k = 1 \dots, d\}$. We then choose $W_{\mathbf{l}} \in \mathcal{A}$ with the largest error indicator $\gamma_{\mathbf{l}}$ and refine that subspace. The subspace is refined by adding all indices $W_{\mathbf{k}}$ with $\mathbf{k} = \mathbf{l} + \mathbf{e}_n$, $n = 1, \dots, d$ that satisfy the following admissibility criterion

$$\mathbf{l} - \mathbf{e}_k \in \mathcal{L} \text{ for } 1 \leq k \leq d, l_k > 1 \quad (5.1)$$

Each subspace that satisfies (5.1) is then placed in the active set $\mathcal{A}$. This process continues until a computational budget limiting the number of model samples (grid points) is reached or a global error indicator drops below a predefined threshold. Pseudo-code for the dimension adaptive algorithm is shown in Algorithm 1.

These `INDICATOR` and `GLOBAL_INDICATOR` routines control which subspaces are added to the sparse grid. The key to adaptive sparse grids working well is greatly influenced by these two routines, which respectively provide accurate estimates of the contribution of a subspace to reducing the error in the interpolant, and the error in the entire interpolant.

Typically indicator functions are functions of the hierarchical surplus values $v_{\mathbf{l},\mathbf{i}}$ of points in the grid. Throughout this paper we will use the indicator

$$\gamma_{\mathbf{l}} = \sum_{\mathbf{i} \in \mathcal{I}_{\mathbf{l}}} |v_{\mathbf{l},\mathbf{i}}| w_{\mathbf{l},\mathbf{i}} \quad (5.2)$$

as the baseline for comparison for the proposed method. Here $w_{\mathbf{l},\mathbf{i}}$ is the quadrature weight of the grid point $\xi_{\mathbf{l},\mathbf{i}}$. In addition, we use the global indicator

$$\eta = \sum_{\mathbf{l} \in \mathcal{A}} \gamma_{\mathbf{l}}$$

Hierarchical surplus indicators such as (5.2) require the model to be evaluated at the associated grid point $\xi_{\mathbf{l},\mathbf{i}}$ before the surplus and thus indicator can be computed. This procedure adds a grid point and then checks if that refinement should have been performed. Efficiency can be gained by using an error indicator that flags the need for refinement without the need to evaluate the simulation model. The a posteriori error estimate (4.1) is one such indicator.

We construct the following error indicators using the a posteriori error estimate for dimension adaptive sparse grids

$$\gamma_1^\varepsilon = \sum_{\mathbf{i} \in \mathcal{I}_1} |\varepsilon(J_{h,n}(\boldsymbol{\xi}_{1,\mathbf{i}}))| w_{1,\mathbf{i}}, \quad \eta^\varepsilon = \sum_{\mathbf{l} \in \mathcal{A}} \gamma_1^\varepsilon \quad (5.3)$$

The use of these indicators requires a minor modification to the interpolate algorithms presented in the literature [13, 19]. The modification, present in Algorithm 1, controls when the model is evaluated at a sparse grid point and when that point is added to the sparse grid. For surplus based refinement the model must be evaluated at the sparse grid when **INDICATOR** is called on a subspace. When using a posteriori refinement the model does not need to be evaluated until **REFINE** is called on the point at which time the point is also added to the sparse grid. This modification means that when using a posteriori refinement each step in the algorithm only requires evaluation of the target function $f(\boldsymbol{\xi}_{1,\mathbf{i}}) \forall \boldsymbol{\xi}_{1,\mathbf{i}} \in W$. In comparison surplus refinement requires evaluation of $f(\boldsymbol{\xi}_{1,\mathbf{i}}) \forall \boldsymbol{\xi}_{1,\mathbf{i}} \in \mathcal{J}$ at each step. Typically the number of points in W will be much less than $\mathcal{J}$ thus a posteriori refinement should be more efficient than surplus based refinement.

Algorithm 1 $\text{INTERP}[f(\boldsymbol{\xi}), \mathcal{L}, \mathcal{A}, \tau, n] \rightarrow f_n, \mathcal{A}, \mathcal{L}$

For a given $\mathcal{L}$ the points in the sparse grid are $\Xi := \bigcup_{\mathbf{l} \in \mathcal{L}} \Xi_1$.

The number of sparse grid points are $N = \#\Xi$

WHILE NOT TERMINATE $[\mathcal{A}, N, \tau, n]$

- $W := \arg \max_{W_1 \in \mathcal{A}} \gamma_1$ % Determine the subspace with the highest priority
 - $\mathcal{A} := \mathcal{A} \setminus W$ % Remove W from the active set
 - IF (NOT using surplus refinement) Evaluate $f(\boldsymbol{\xi}_{1,\mathbf{i}}) \forall \boldsymbol{\xi}_{1,\mathbf{i}} \in W$
 - $\mathcal{J} := \text{REFINE}[W, \mathcal{L}]$ % Find all admissible forward neighbors of W
 - IF (using surplus refinement) Evaluate $f(\boldsymbol{\xi}_{1,\mathbf{i}}) \forall \boldsymbol{\xi}_{1,\mathbf{i}} \in \mathcal{J}$
 - $\gamma_1 := \text{INDICATOR}[W_1] \forall W_1 \in \mathcal{J}$ % Calculate the priority of the neighbors
 - $\mathcal{A} := \mathcal{A} \cup \mathcal{J}$ % Add the forward neighbors to the active index set
-

Algorithm 2 TERMINATE[$\mathcal{A}, N, \tau, n$]

- $\eta := \text{GLOBAL_INDICATOR}[\mathcal{A}]$
 - IF $\mathcal{A} = \emptyset$ or $N \geq n$ or $\eta < \tau$ RETURN TRUE
 - ELSE RETURN FALSE
-

5.2. Local adaptivity

Locally-adapted sparse grids can be adapted using a conceptually similar method to that employed for dimension adapted grids. Instead of refining the grid subspace by subspace, locally-adaptive grids are adapted point by point. When a new point is chosen for refinement it is added to the sparse grid. Children of this point are then found and added to an active set $\mathcal{A}_{\text{local}}$. In this paper we will consider the typical local refinement strategy [20] and the simultaneous dimension and local refinement proposed in [19] termed generalized local refinement. In this manuscript we will refer to the former as traditional local refinement and the later as generalized local refinement. Error indicators, are then computed at each point in $\mathcal{A}_{\text{local}}$. The next point chosen for refinement is simply the point $\xi_{\mathbf{l}, \mathbf{i}}$ in $\mathcal{A}_{\text{local}}$ with the largest error indicator $\gamma_{\mathbf{l}, \mathbf{i}}$. The local adaptive sparse grid procedure can be obtained from Algorithm 1 by defining $\mathcal{A}$ and $\mathcal{L}$ to contain points rather than subspaces and W and W_1 to be grid points not subspaces. The **REFINE** routine must also be changed as discussed in [19].

Throughout this paper we will use the indicators

$$\gamma_{\mathbf{l}, \mathbf{i}} = |v_{\mathbf{l}, \mathbf{i}}| w_{\mathbf{l}, \mathbf{i}}, \quad \eta = \sum_{\{\mathbf{l}, \mathbf{i}\} \in \mathcal{A}} \gamma_{\mathbf{l}, \mathbf{i}} \quad (5.4)$$

as the baseline for comparison for the proposed method. Again we can use a posteriori error estimates to construct the alternative refinement metric

$$\gamma_{\mathbf{l}, \mathbf{i}}^\varepsilon = |\varepsilon(J_{h,n}(\xi_{\mathbf{l}, \mathbf{i}}))| w_{\mathbf{l}, \mathbf{i}}, \quad \eta^\varepsilon = \sum_{\{\mathbf{l}, \mathbf{i}\} \in \mathcal{A}} \gamma_{\mathbf{l}, \mathbf{i}}^\varepsilon \quad (5.5)$$

We remark that when using local refinement one should ensure all the ancestors of a point in the sparse grid have previously been added to the sparse grid. Missing ancestors must be added recursively before a child is added. The use of a posteriori error guided refinement affects when the **INTERP** algorithm adds ancestors to the sparse grid. When employing a posteriori based refinement, function values at ancestor points are not necessary to compute

the error estimate at a child point. One has the choice to add missing ancestors either when the error estimate of a child is computed or later when the child point is added to the sparse grid. We found during our investigation that adding ancestors before computing the error estimate produced a more accurate error estimate at the candidate points which resulted in a more accurate sparse grid for a fixed number of sample functions. These results are not shown as the authors believed this minor point would distract from the main conclusions that can be drawn from the results section.

Throughout the remainder of this manuscript we will use the following notation to denote the various types of refinement strategies. Let γ_1 and γ_1^ε respectively denote, hierarchical surplus and a posteriori based dimension based refinement. Let hierarchical surplus based traditional local refinement [20] and generalized local refinement [19] be denoted by $\gamma_{1,i}^{\text{trad}}$ and $\gamma_{1,i}^{\text{gen}}$, respectively. Finally we define $\gamma_{1,i}^{\varepsilon,\text{gen}}$ to be a posteriori error based local generalized refinement.

6. Sparse grid approximations of the error

As previously mentioned, the evaluation of the a posteriori error estimate using (4.1) requires the full forward and adjoint approximations at $\xi \in \Gamma_\xi$ as well as the evaluation of the space-time weighted-residual. While this is usually much cheaper than solving the forward problem at ξ , the computation cost in producing the error estimate at ξ should not be neglected. If we only need to compute a small number of samples, then this is probably not a significant issue. However, we often require a very large number of samples of the surrogate, and, in turn, a very large number of error estimates. To mitigate this issue, we propose constructing a surrogate of the error estimate. This can significantly reduce the number of error estimates required and the surrogate of the error estimate can be chosen to have similar accuracy as the point-wise evaluation of (4.1).

Given a sparse grid approximation of the quantity of interest $J_{h,n}(\xi)$ we can use a finite number of error estimates $\varepsilon(J_{h,n}(\xi_{1,i}))$ at new grid points to build an enhanced sparse grid. Once built the enhanced sparse grid interpolant can be sampled directly during post-processing, thus removing any further need to evaluate (4.1).

The procedure needed to construct an enhanced sparse grid is outlined in Algorithm 3. The algorithm consist of two steps. The first step is used to construct sparse grid approximations of $J_h(\xi)$, forward solution $\mathbf{z}_h$ and the adjoint solution ϕ_h . The second step uses these sparse grids as the basis of a sparse grid approximation of $J_n^\varepsilon(\xi)$, Specifically the final $\mathcal{A}$ and $\mathcal{L}$ of

$J_{h,n}(\boldsymbol{\xi})$ are used as the initial starting point when building $J_{n,m}^\varepsilon$ and are then evolved by sampling $J_n^\varepsilon(\boldsymbol{\xi})$ at new sparse grid points. We emphasize that the second step no longer involves full PDE solves, but only evaluation of the sparse grids approximation of the PDE solutions and a weighted residual needed to compute $\varepsilon(\mathbf{z}_{h,n}, \boldsymbol{\phi}_{h,n})$.

Algorithm 3 ENHANCED-INTERP[$J_h(\boldsymbol{\xi}), \mathbf{z}_h(\boldsymbol{\xi}), \boldsymbol{\phi}_h(\boldsymbol{\xi}), \mathcal{L}, \mathcal{A}, \tau, n$] $\rightarrow J_n^\varepsilon$

Define C to be the ratio of the cost of evaluating the error estimate to the combined cost of a forward and adjoint equation evaluation

- $J_h(\boldsymbol{\xi})_{h,n}, \mathcal{A}_n, \mathcal{L}_n = \text{INTERP}[J_h(\boldsymbol{\xi}), \mathcal{A}, \mathcal{L}, \tau, n/4]$
- $\mathbf{z}_{h,n}(\boldsymbol{\xi}, \mathbf{x}, \mathbf{t}), \mathcal{A}_n, \mathcal{L}_n = \text{INTERP}[\mathbf{z}_h(\boldsymbol{\xi}, \mathbf{x}, \mathbf{t}), \mathcal{A}, \mathcal{L}, \tau, n/4]$
- $\boldsymbol{\phi}_{h,n}(\boldsymbol{\xi}, \mathbf{x}, \mathbf{t}), \mathcal{A}_n, \mathcal{L}_n = \text{INTERP}[\boldsymbol{\phi}_h(\boldsymbol{\xi}, \mathbf{x}, \mathbf{t}), \mathcal{A}, \mathcal{L}, \tau, n/4]$
- $\delta_{\max} = \max_{\boldsymbol{\xi}_{1,i} \in \mathcal{A}} \delta(J_h(\boldsymbol{\xi}_{1,i}))$
- $\gamma_{\max} = \max_{\boldsymbol{\xi}_{1,i} \in \mathcal{A}} \gamma_{1,i}$
- $\tau_\varepsilon = \max(\delta_{\max}, \gamma_{\max}^2)$
- $J_{n,m}^\varepsilon(\boldsymbol{\xi}), \mathcal{A}_{n,m}, \mathcal{L}_{n,m} = \text{INTERP}[J_{h,n}(\boldsymbol{\xi}) + \varepsilon(\mathbf{z}_{h,n}, \boldsymbol{\phi}_{h,n}), \mathcal{A}_n, \mathcal{L}_n, \tau_\varepsilon, Cn/2]$

% $\mathbf{z}_{h,n}$ and $\boldsymbol{\phi}_{h,n}$ are needed to compute $\varepsilon(\mathbf{z}_{h,n}, \boldsymbol{\phi}_{h,n})$.

6.1. Adaptivity

When using dimension-adaptive sparse grids we begin with an initial grid with subspaces in $\mathcal{L}_n$ and a set of candidate subspaces $\mathcal{A}_n$ generated when building the sparse grid approximation of the quantity of interest, forward solution and the adjoint solution. When using surplus refinement $\mathcal{A}_n \subset \mathcal{L}_n$ and when using a posteriori error refinement $\mathcal{L}_n \cap \mathcal{A}_n = \emptyset$. This difference arises from the fact that surplus refinement requires the model to be evaluated at all the points in the candidate subspaces where as a posteriori refinement does not. The set $\mathcal{L}_n$ at this stage of the algorithm are the gray boxes in Figure 1.

Given the sets $\mathcal{A}_n$ and $\mathcal{L}_n$ we simply use $\text{INTERP}[J_n^\varepsilon(\boldsymbol{\xi}), \mathcal{L}_n, \mathcal{A}_n, \tau_\varepsilon, m]$ to construct the enhanced sparse grid $J_{n,m}^\varepsilon := I_m[J_n^\varepsilon(\boldsymbol{\xi})]$. When the algorithm terminates $\mathcal{L}$ will include the red and gray boxes in Figure 1 and $\mathcal{A} = \emptyset$.

Aside from noting this minor distinction, constructing $J_{n,m}^\varepsilon$ requires one additional step. Before $\text{INTERP}[J_n^\varepsilon(\boldsymbol{\xi}), \mathcal{L}_n, \mathcal{A}_n, \tau_\varepsilon, m]$ is called the function

Figure 1: The two phases undertaken when building an enhanced dimension adaptive sparse grid interpolant J_n^ε . The first phase consist of building the un-enhanced approximation $J_{h,n}(\xi)$. The gray boxes on the left represent the 2D sparse grid subspaces $\mathbf{l} = (l_1, l_2)$ used to build $J_{h,n}(\xi)$ and the black points on the right represent the corresponding points in that sparse grid. At each of the black points both the forward and adjoint equations are solved. The second phase involves continuing refinement until the sparse grid adequately resolves the enhanced function J_n^ε . The red boxes represent the additional subspaces used to build the enhanced approximation $J_{h,n}^\varepsilon$. The red points are the associated additional points that are added to the sparse grid. Unlike the black points the red points only require evaluation of the residual (to compute the a posteriori error estimate) and the cheaply evaluated approximation $J_{h,n}(\xi)$. The red boxes are not to be confused with active indices which cannot have the structure shown in this figure.

value at all points in $\mathcal{L}_n$ must be updated to include the physical discretization error estimate $\delta(J_h(\xi))$. That is for all $\xi_{1,i} \in \mathcal{L}_n$ set the function value at that point to $J_h(\xi_{1,i}) + \delta(J_h(\xi_{1,i}))$. The hierarchical surpluses at each of these points must also be updated accordingly. We do this because any new points to the enhanced sparse grid will have function values $J_{h,n}(\xi_{1,i}) + \varepsilon(J_{h,n}(\xi_{1,i}))$, where $\varepsilon(J_{h,n}(\xi_{1,i}))$ includes the deterministic error estimate. We again remark that in order to compute the physical discretization error estimate, the physical discretization used to compute the adjoint solution must be higher-order than the method used to compute the forward solution. This is standard practice in adjoint-based error estimation for deterministic problems.

The method used to construct an enhanced sparse grid using local refinement is similar to that employed when using dimension refinement. Again we must simply define the initial sets $\mathcal{L}_n$ and $\mathcal{A}_n$ and then call INTERP after adjusting the function values and hierarchical surpluses for all $\xi_{1,i} \in \mathcal{L}_n$.

6.2. Balancing the stochastic and physical discretization errors

When quantifying the uncertainty of model predictions, both the deterministic and stochastic discretization errors must be accounted for. It

is inefficient to reduce the stochastic error to a level below the error introduced by the physical discretization. In the following, we assume that physical discretization, and thus the error induced by the physical discretization, is fixed. In order to balance the stochastic error with this fixed physical discretization error, we must be able to handle two regimes: first that our computational budget is large enough to drive the stochastic error below the physical discretization error; and secondly the computational budget is insufficient to force the stochastic error below the physical discretization error.

In the first regime we must be able to identify when the sparse grid is sufficiently refined. We use τ_ε to terminate refinement of $J_{n,m}^\varepsilon$ when the interpolation error is approximately equal to the physical discretization error. Specifically, we set τ_ε to be the maximum of the approximate physical discretization error $\delta_{\max}$ and the approximate stochastic error $\gamma_{\max}$ in the enhanced sparse grid. The error in the error estimate is the product of the error in the approximation of the forward solution and the error in the approximation of the adjoint solution [7, 6, 8]. Consequently a reasonable approximation of the potential accuracy of the enhanced sparse grid is the square of the maximum indicator $\gamma_{1,i}$ of all points in $\mathcal{A}$. That is we set $\tau_\varepsilon = \eta^2$

In the second regime, the goal is to reduce the total error as much as possible within a computational budget of n computational units. To do this we must consider the costs of solving the forward problem, solving the adjoint problem, and evaluating the error estimate (4.1). From our experiments we found it to be advantageous to limit the number of forward and adjoint solves to be 1/2 of the computational budget, and to utilize the remaining effort to evaluate $\varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})$ to build the enhanced sparse grid. Here we have made the reasonable assumption that the cost of one forward solve and one adjoint solve are approximately equal, thus the choice of $n/4$ when constructing $\mathbf{z}_{h,n/4}$ and $\phi_{h,n/4}$. The cost in solving the adjoint problem may actually be smaller than the cost in solving the forward problem, even if a higher order method is used for the adjoint, since the adjoint problem is always linear.

We remark that these termination conditions removes the need for knowledge of what computational budget regime is active.

7. Results

In this section, we provide several numerical examples to illustrate the proposed methodology. In all of the figures presented, the expression left

of the colon in the legend denotes the quantity being approximated and the expression on the right denotes the type of refinement employed. For definitions of the refinement types refer to the end of Section 5.

7.1. Heterogeneous diffusion equation

In this section, we investigate the performance of the proposed methodology when applied to a Poisson equation with an uncertain heterogeneous diffusion coefficient. Attention is restricted to the one-dimensional physical space to avoid unnecessary complexity. The procedure described here can easily be extended to higher physical dimensions.

Consider the following problem with $d \geq 1$ random dimensions:

$$-\frac{d}{dx} \left[a(x, \boldsymbol{\xi}) \frac{d\mathbf{z}}{dx}(x, \boldsymbol{\xi}) \right] = 10, \quad (x, \boldsymbol{\xi}) \in (0, 1) \times I_{\boldsymbol{\xi}} \quad (7.1)$$

subject to the physical boundary conditions

$$\mathbf{z}(0, \boldsymbol{\xi}) = 0, \quad \mathbf{z}(1, \boldsymbol{\xi}) = 0. \quad (7.2)$$

Furthermore assume that the log random diffusivity satisfies

$$a(x, \boldsymbol{\xi}) = \bar{a} + \sigma_a \sum_{k=1}^d \sqrt{\lambda_k} \phi_k(x) \xi_k \quad (7.3)$$

where $\{\lambda_k\}_{k=1}^d$ and $\{\phi_k(x)\}_{k=1}^d$ are, respectively, the eigenvalues and eigenfunctions of the covariance kernel

$$C_a(x_1, x_2) = \exp \left[-\frac{(x_1 - x_2)^2}{2l_c} \right]. \quad (7.4)$$

The variability of the diffusivity field (7.3) is controlled by σ_a and the correlation length l_c which determines the decay of the eigenvalues λ_k . In the following we set $d = 25$, $\sigma_a = 1.$, $l_c = 0.1$, $\bar{a} = 0$ and define $\xi_k \in [-1, 1]$, $k = 1, \dots, d$ to be independent and identically distributed uniformly random variables.

We are interested in quantifying the uncertainty in the solution $\mathbf{z}$ at $x = 0.5$ caused by the aforementioned random diffusivity field (7.3). Accordingly, we define our quantity of interest to be the linear functional

$$J(\boldsymbol{\xi}) = (\psi(x), \mathbf{z}(x, \boldsymbol{\xi})) , \quad \psi(x) = C_s \exp(-100(x - .5)^2),$$

where C_s is a scaling constant chosen so that $\int_0^1 \psi(x) dx = 1$. The adjoint problem corresponding to this quantity of interest is given by

$$-\frac{d}{dx} \left[a(x, \boldsymbol{\xi}) \frac{d\phi}{dx}(x, \boldsymbol{\xi}) \right] = \psi(x), \quad (x, \boldsymbol{\xi}) \in (0, 1) \times I_{\boldsymbol{\xi}} \quad (7.5)$$

subject to the physical boundary conditions

$$\phi(0, \boldsymbol{\xi}) = 0, \quad \phi(1, \boldsymbol{\xi}) = 0. \quad (7.6)$$

We approximate the solution to (7.1) using a standard continuous Galerkin Finite Element Method (FEM) on a uniform mesh with 100 elements ($h = 0.01$) and piecewise linear basis functions. Similarly, to solve (7.5) we use a continuous Galerkin method with piecewise quadratic basis functions. A higher-order discretization is used to solve the adjoint problem so that estimates of the physical discretization error (due to the finite element approximation) can be obtained.

7.1.1. Adjoint based estimates of the stochastic error

Here, we demonstrate that sparse grid estimates of the linear functional $J_{h,n}(\boldsymbol{\xi})$ can be made significantly more accurate by combining these estimates with a posteriori error estimates $\varepsilon(\mathbf{z}_{h,n}, \boldsymbol{\phi}_{h,n})$.

Figure 2 compares the errors $\|J_h(\boldsymbol{\xi}) - J_{h,n}(\boldsymbol{\xi})\|_{\ell_2(\Gamma_{\boldsymbol{\xi}})}$ of un-enhanced sparse surplus-based, dimension-adaptive, Clenshaw-Curtis grids with the errors of sparse grid approximations enhanced with a posterior error estimates $\|J_h(\boldsymbol{\xi}) - (J_{h,n}(\boldsymbol{\xi}) + \varepsilon(\mathbf{z}_{h,n}, \boldsymbol{\phi}_{h,n}))\|_{\ell_2(\Gamma_{\boldsymbol{\xi}})}$. Here and throughout the paper discrete ℓ_2 errors are computed using 100,000 Latin hypercube samples of $\Gamma_{\boldsymbol{\xi}}$. We observe that the rate of convergence of the un-enhanced sparse grid is ≈ 1.8 and the rate for the enhanced grid is ≈ 3.3 . This is consistent with the theoretically results in Section 4 which predicts the rate of the enhanced sparse grid to be approximately twice the rate of the un-enhanced sparse grid.

Convergence of errors is shown with respect to increasing computational cost. Here, and throughout this paper, we define one unit of cost to be the computational cost of one solve of the forward model and assume that the cost of solving the adjoint equation is also equal to one unit. For this example, we ignore the cost of computing the error estimate. We will challenge this last assumption shortly.

7.1.2. Balancing physical and stochastic discretization errors

When quantifying the uncertainty of model predictions, both the deterministic and stochastic discretization errors must be accounted for. It is inefficient to reduce the stochastic error to a level below the error introduced

by the physical discretization. This statement is supported by Figure 2. When the stochastic error is smaller than the physical discretization error the accuracy in the enhanced sparse grid stagnates at approximately the physical discretization error.

Figure 2: Errors, in the Clenshaw-Curtis, dimension-adaptive, un-enhanced and enhanced sparse grid approximations of (7.4) for the 25d heat equation.

7.1.3. Reducing the cost of computing a posterior error estimates

Although typically ignored in the literature, the cost of computing $\varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})$ is non-trivial. For this example the cost of evaluating $\varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})$ is approximately 1/25 of a computational unit and the cost of evaluating $\varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})$ at 100,000 points, needed to compute the ℓ_2 error, often outweighs the cost of the forward and adjoint solves.

To reduce the cost of producing many a posterior error estimates we build an approximation $J_{n,m}^\varepsilon$ of $J_{h,n}(\xi) + \varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})$ using the approach outlined in Algorithm 3. Figure 3 demonstrates that the performance of Algorithm 3 is dependent on the parameters τ and n . Simply driving the sparse grid error indicator to zero is inefficient. After a certain number of samples, adding more points (each requiring a a posteriori error estimate)

to the enhanced sparse grid has no effect on reducing the ℓ_2 error of the approximation. However if τ and n are set appropriately as described in 6.2 the number of ‘wasted’ samples can be minimized.

Figure 3: The decay of the ℓ_2 error vs the decay in the error indicator for the enhanced sparse grid $J_{n,m}^\varepsilon$, starting from an initial grid with 51 points (left) and 103 points (right).

Figure 4 plots the convergence of the un-enhanced surplus-based dimension-adaptive Clenshaw-Curtis sparse grids $J_{h,n}(\boldsymbol{\xi})$, J_n^ε and $J_{n,m}^\varepsilon$. The enhanced sparse grid $J_{n,m}^\varepsilon$ is significantly more accurate than the un-enhanced grid, however $J_{n,m}^\varepsilon$ is not as accurate as J_n^ε . Building $J_{n,m}^\varepsilon$ reduces the number of samples needed to generate enhanced sparse grid samples at the expense of a slight degradation in accuracy. This degradation is acceptable as in practice, adding $\varepsilon(\mathbf{z}_{h,n}, \boldsymbol{\phi}_{h,n})$ to J_n^ε is typically computationally infeasible². From Figure 4 we also observe that our choice of τ and n stops refinement of $J_{n,m}^\varepsilon$ before the stochastic error is forced below the deterministic error $\delta_{\max} = 8.41 \times 10^{-5}$.

Figure 5 plots the convergence of the un-enhanced surplus-based locally-adaptive sparse grids $J_{h,n}(\boldsymbol{\xi})$, J_n^ε and $J_{n,m}^\varepsilon$. The same conclusions drawn when using dimension adaptivity can also be made here. The improvement obtained by using $J_{n,m}^\varepsilon$ instead of $J_{h,n}(\boldsymbol{\xi})$ is reduced however this can be addressed by the use of different refinement strategies discussed in Section 7.3

7.2. Non-linear coupled system of ODEs

For our second example, consider the non-linear system of ordinary differential equations governing a competitive Lotka–Volterra model of the

²The computational cost of $J_{n,m}^\varepsilon$ shown in Figure 4 includes the cost of evaluating the residual to compute the error estimates needed at the sparse grid points, whereas we have assumed that computing $\varepsilon(\mathbf{z}_{h,n}, \boldsymbol{\phi}_{h,n})$ for J_n^ε has no cost.

Figure 4: Comparison of the ℓ_2 convergence of the enhanced dimension-adaptive sparse grid using direct evaluation of the error estimate J_n^ε and a sparse grid approximation of the enhanced QOI $J_{n,m}^\varepsilon$.

population dynamics of species competing for some common resource. The model is given by

$$\begin{cases} \frac{d\mathbf{z}_i}{dt} = r_i \mathbf{z}_i \left(1 - \sum_{j=1}^3 \alpha_{ij} \mathbf{z}_j \right), & t \in (0, 10], \\ \mathbf{z}_i(0) = \mathbf{z}_{i,0} \end{cases}, \quad (7.7)$$

for $i = 1, 2, 3$. The initial condition, $\mathbf{z}_{i,0}$, and the self-interacting terms, α_{ii} , are given, but the remaining interaction parameters, α_{ij} with $i \neq j$ as well as the re-productivity parameters, r_i , are unknown. We assume that these parameters are each uniformly distributed on $[0.3, 0.7]$. We approximate the solution to (7.7) in time using a Backward Euler method with 1000 time steps ($\Delta t = 0.01$) which gives a deterministic error of approximately 1.00×10^{-4} .

The quantity of interest is the population of the third species at the final

Figure 5: Comparison of the ℓ_2 convergence of the enhanced locally-adaptive sparse grid using direct evaluation of the error estimate and a sparse grid approximation of the enhanced QOI

time, $\mathbf{z}_3(10)$. The corresponding adjoint problem is

$$\begin{cases} -\frac{d\phi_i}{dt} = r_i\phi_i \left(1 - \sum_{j=1}^3 \alpha_{ij}\mathbf{z}_j\right) + r_i\mathbf{z}_i \left(1 - \sum_{j=1}^3 \alpha_{ji}\phi_j\right), & t \in (10, 0], \\ \phi_i(10) = 0, \\ \phi_3(10) = 1. \end{cases} \quad i = 1, 2, \quad (7.8)$$

We approximate the adjoint solution in time using a second-order Crank-Nicholson method with the same number of time steps.

Figure 6 and Figure 7 compare the convergence of $J_{h,n}(\boldsymbol{\xi})$, J_n^ε and $J_{n,m}^\varepsilon$ when using a surplus based dimension-adaptive Clenshaw-Curtis sparse grid and surplus locally adaptive sparse grid, respectively. Again, significant increases in accuracy can be achieved by using $J_{n,m}^\varepsilon$ instead of $J_{h,n}(\boldsymbol{\xi})$. In addition we again see that we are able to correctly stop refinement approximately when the stochastic error of the sparse grid approximation is the same order as the physical discretization error. We emphasize that, as in the previous example, the accuracy of $J_{n,m}^\varepsilon$ can be forced towards the accuracy

of the best possible enhanced approximation $J_n^\varepsilon + \varepsilon(\mathbf{z}_{h,n}, \phi_{h,n})$ at the cost of additional residual calculations (error estimates) at new sparse grid points.

Figure 6: Comparison of the ℓ_2 convergence of the enhanced dimension-adaptive sparse grid using direct evaluation of the error estimate and a sparse grid approximation of the enhanced QOI

7.3. *A posteriori* error based refinement strategies

In this section we demonstrate the utility of using a posteriori error based refinement strategies outlined in Section 5. Specifically we consider a posteriori refinement indicators for both dimension and locally adapted sparse grid. In the case of local refinement we also consider the simultaneous dimension and local refinement in addition to the typical local refinement strategy used in all examples thus far. We remark that in all the figures shown in the remainder of this paper, the cost of computing the sparse grid approximation, shown on the horizontal axis, includes the cost of evaluating the a posteriori error estimate, building the enhanced the sparse grid approximation, and calculating the refinement indicators.

Figure 7: Comparison of the ℓ_2 convergence of the enhanced locally-adaptive sparse grid using direct evaluation of the error estimate and a sparse grid approximation of the enhanced QOI

7.3.1. Dimension refinement

Figure 8 compares refinement strategies for the 25d heat equation problem presented in Section 7.1. It is clear that the use of a posteriori refinement results in significant increases in efficiency. This result is due to the fact that the a posteriori subspace refinement indicator γ_1 only requires computation of error estimates which are relatively cheap compared to the evaluation of the forward model. The function evaluations that are needed by surplus refinement to probe candidate subspaces are often redundant and this redundant evaluation can be avoided by using a posteriori refinement.

The high-dimensional nature and the strong degree of anisotropy of this problem are an ideal setting for a posteriori refinement. If the function was lower dimensional or have less anisotropy then surplus refinement would be more competitive as the amount of redundant model evaluations would decrease. This point is supported by Figure 9 which compares a posteriori and surplus dimension adaptive sparse grids when used to solve the 9d predator-prey model presented in Section 7.2. The predator prey model is

less anisotropic than the diffusion model and thus the benefits of a-posteriori refinement are reduced, although the benefits are still positive.

Figure 8: Comparison of dimension-adaptive refinement strategies when applied to the 25d heat equation

7.3.2. Local refinement

A posteriori refinement can also be used in conjunction with local refinement. Figure 10 compares refinement strategies for the 25d heat equation. The use of generalized local refinement results in a vast improvement over traditional local refinement. The best refinement strategy for this problem is to use a posteriori refinement indicators with the generalized local refinement. This result is due to the high-dimension of the problem and the high degree of anisotropy.

Figure 11 compares refinement strategies for the 9d predator prey model. Due to the low degree of anisotropy of this function generalized local refinement provides no improvement over traditional local refinement. In the case of the surplus-based enhanced sparse grid the accuracy is slightly worse. Traditional surplus based refinement is the most accurate strategy here,

Figure 9: Comparison of dimension-adaptive refinement strategies when applied to 9d predator prey model

but the a posteriori error based generalized local refinement has comparable accuracy.

The best refinement strategy is problem dependent. The authors recommend basing the choice of refinement strategy by performing studies such as the one presented in [25]. However in the absence of such a investigation and based upon the examples presented in this paper, the authors recommend using a posteriori based refinement for both dimension adaptive and local adaptive sparse grids. When using local adaptivity a posteriori based refinement should be combined with the generalized local adaptive procedure presented in [19].

Finally we remark that the primary concern of this section was to discuss the efficacy of a posteriori and surplus based refinement and not the strengths and weaknesses of dimension vs local refinement. Examples using both dimension and local refinement were provided solely to demonstrate that either type of refinement can be used in conjunction with a posteriori error estimates. The benefit of local and dimension refinement is dependent on the regularity of the PDE solutions. Practically determining which

Figure 10: Comparison of local refinement strategies when applied to the 25d heat equation

approach is better is an open question.

8. Conclusions

In this paper we present an algorithm for adaptive sparse grid approximations of quantities of interest computed from discretized partial differential equations. We use adjoint-based a posteriori error estimates of the interpolation error in the sparse grid to enhance the sparse grid approximation. Using a number of numerical examples we show that utilizing a posteriori error estimates provides significantly more accurate functional values for random samples of the sparse grid approximation. The cost of computing these error estimates can be non-trivial and thus we provide a practical method for enhancing a sparse grid approximation using only a finite set of error estimates.

Aside from using a posteriori error estimates to enhance an approximation we also demonstrate that error estimates can be used to increase the efficiency of adaptive sampling. Traditional sparse grid adaptivity employs

Figure 11: Comparison of local refinement strategies when applied to the 9d predator prey model

error indicators based upon the hierarchical surplus are used to flag dimensions or local regions for refinement. However such approaches require the model to be evaluated at a new point before one can determine if refinement should have taken place. In this paper we numerically demonstrated that refinement using a posteriori error estimates can significantly reduce the amount of redundant sampling compared when compared to traditional hierarchical refinement. Refinement indicators based upon a posteriori error estimates are most effective when the quantity of interest being approximated is anisotropic, however even in the absence of anisotropy a posteriori indicators perform at least comparably to traditional hierarchical surplus based indicators.

In combination with the aforementioned enhancement and refinement procedures we use a posteriori error estimates to ensure that the sparse grid is not refined beyond the point at which the stochastic interpolation error is below the physical discretization error. The methodology presented provides a practical means of balancing the stochastic and deterministic discretization errors.

References

- [1] BABUSKA, I., NOBILE, F., AND TEMPONE, R. A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM Journal on Numerical Analysis* 45, 3 (2007), 1005–1034.
- [2] BANGERTH, W., AND RANNACHER, R. *Adaptive Finite Element Methods for Differential Equations*. Birkhauser Verlag, 2003.
- [3] BECKER, R., AND RANNACHER, R. An optimal control approach to a posteriori error estimation in finite element methods. *Acta Numerica* 10 (2001), 1–102.
- [4] BRENNER, S. C., AND SCOTT, L. R. *The Mathematical Theory of Finite Element Methods*. Springer, 2002.
- [5] BUNGARTZ, H.-J., AND GRIEBEL, M. Sparse grids. *Acta Numerica* 13 (2004), 147–269.
- [6] BUTLER, T., CONSTANTINE, P., AND WILDEY, T. A posteriori error analysis of parameterized linear systems using spectral methods. *SIAM J. Matrix Anal. Appl.* 33 (2012), 195–209.
- [7] BUTLER, T., DAWSON, C., AND WILDEY, T. A posteriori error analysis of stochastic spectral methods. *SIAM J. Sci. Comput.* 33 (2011), 1267–1291.
- [8] BUTLER, T., DAWSON, C., AND WILDEY, T. Propagation of uncertainties using improved surrogate models. *SIAM/ASA Journal on Uncertainty Quantification* 1, 1 (2013), 164–191.
- [9] CIARLET, P. G. Basic error estimates for elliptic problems. In *Handbook of numerical analysis, Vol. II*. North-Holland, Amsterdam, 1991, pp. 17–351.
- [10] ERIKSSON, K., ESTEP, D., HANSBO, P., AND JOHNSON, C. *Computational differential equations*. Cambridge University Press, Cambridge, 1996.
- [11] ESTEP, D., LARSON, M. G., AND WILLIAMS, R. D. Estimating the error of numerical solutions of systems of reaction-diffusion equations. *Mem. Amer. Math. Soc.* 146, 696 (2000), viii+109.

- [12] ESTEP, D., TAVENER, S., AND WILDEY, T. A posteriori error estimation and adaptive mesh refinement for a multiscale operator decomposition approach to fluid-solid heat transfer. *J. Comput. Phys.* 229 (June 2010), 4143–4158.
- [13] GERSTNER, T., AND GRIEBEL, M. Dimension-adaptive tensor-product quadrature. *Computing* 71, 1 (SEP 2003), 65–87.
- [14] GHANEM, R., AND SPANOS, P. *Stochastic Finite Elements: A Spectral Approach*. Springer Verlag, New York, 2002.
- [15] GILES, M. B., AND SÜLI, E. Adjoint methods for pdes: a posteriori error analysis and postprocessing by duality. *Acta Numerica* 11 (1 2002), 145–236.
- [16] GRIEBEL, M. Adaptive sparse grid multilevel methods for elliptic PDEs based on finite differences. *Computing* 61, 2 (1998), 151–179.
- [17] GUNZBURGER, M. D. *Finite Element Methods for Viscous Incompressible Flows*. Academic Press, 1989.
- [18] JAKEMAN, J., ELDRED, M., AND XIU, D. Numerical approach for quantification of epistemic uncertainty. *J. Comput. Phys.* 229, 12 (2010), 4648–4663.
- [19] JAKEMAN, J., AND ROBERTS, S. Local and dimension adaptive stochastic collocation for uncertainty quantification. In *Sparse Grids and Applications*, J. Garcke and M. Griebel, Eds., vol. 88 of *Lecture Notes in Computational Science and Engineering*. Springer Berlin Heidelberg, 2013, pp. 181–203.
- [20] MA, X., AND ZABARAS, N. An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations. *J. Comput. Phys.* 228 (2009), 3084–3113.
- [21] NOBILE, F., TEMPONE, R., AND WEBSTER, C. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.* 46, 5 (2008), 2309–2345.
- [22] NOBILE, F., TEMPONE, R., AND WEBSTER, C. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis* 46, 5 (2008), 2309–2345.

- [23] ODEN, J. T., AND PRUDHOMME, S. Goal-oriented error estimation and adaptivity for the finite element method. *Computers & mathematics with applications* 41, 5 (2001), 735–756.
- [24] RASMUSSEN, C. E. Gaussian processes for machine learning.
- [25] WINOKUR, J., CONRAD, P., SRAJ, I., KNIO, O., SRINIVASAN, A., THACKER, W., MARZOUK, Y., AND ISKANDARANI, M. A priori testing of sparse adaptive polynomial chaos expansions using an ocean general circulation model database. *Computational Geosciences* (2013), 1–13.
- [26] XIU, D., AND KARNIADAKIS, G. The Wiener-Askey Polynomial Chaos for stochastic differential equations. *SIAM J. Sci. Comput.* 24, 2 (2002), 619–644.

active-indices-example.png