

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

Jakeman, J.D.^{a,1,*}, Willey, 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. We show that 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

1. Introduction

Partial Differential Equations (PDE) are used to simulate a wide range of phenomenon and are often used to inform design decisions and to esti-

*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.

mate risk in systems with large human and/or financial impact but with limited capacity for experimentation. Given the importance of these applications 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 these governing equations. Many applications 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 [12, 24], sparse grid interpolation [17, 18] and Gaussian Process (GP) models [21]. 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 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

$$\|J(\boldsymbol{\xi}) - J_{h,n}(\boldsymbol{\xi})\| \leq \underbrace{\|J(\boldsymbol{\xi}) - J_h(\boldsymbol{\xi})\|}_I + \underbrace{\|J_h(\boldsymbol{\xi}) - J_{h,n}(\boldsymbol{\xi})\|}_{II} \quad (1.1)$$

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) [1, 8, 13, 20], 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 [5] and further analyzed in [4, 6], 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 [4, 6], 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 [19]. 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

- Show that the enhancement results in [4] also extend to adaptive sparse grid approximations.
- Present new adaptivity strategies for sparse grids based on a posteriori error estimates.
- Demonstrate that, for a given computational budget, better accuracy can be obtained using the proposed approach.
- Provide a stopping criterion, based on the physical discretization error, to avoid over-resolving the sparse grid.

The remainder of this paper is organized as follows. Section 2 introduces a general model problem. In Section 3 we recall the standard adjoint-based posteriori error analysis for deterministic PDEs. Section 4 presents the standard adaptive sparse grid algorithms. In Section 5 we derive an a posteriori error estimate for samples of a sparse grid surrogate. In Section 6, we present the new adaptive strategies using the error estimates. Section 7 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 8 and our conclusions are presented in Section 9.

2. General nonlinear problem and notation

We consider the following system of partial differential equations,

$$\frac{\partial \mathbf{z}}{\partial t} + \mathbf{A}(\boldsymbol{\xi}; \mathbf{z}) = \mathbf{0}, \quad (2.1)$$

defined on $\Omega \times (0, T]$ where $\Omega \subset \mathbb{R}^s$, $s = 2, 3$, 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 $\mathbf{z} := \mathbf{z}(x, t, \boldsymbol{\xi})$ is also uncertain and may be modeled as a random process for which we will construct a surrogate approximation. 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 (2.1) is well-posed in the sense that there exists a solution for a. e. $\boldsymbol{\xi} \in \Gamma_{\boldsymbol{\xi}}$.

We use $(\cdot, \cdot)_{\mathcal{D}}$ to 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} . We let \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 frequently use $\mathbf{H}^m(\Omega)$ to denote the obvious generalization to vector valued functions.

We assume that (2.1) has an equivalent variational formulation seeking, for any $\boldsymbol{\xi} \in \Gamma_{\boldsymbol{\xi}}$, $\mathbf{z} \in \mathbf{V} \times (0, T]$ such that,

$$\int_0^T \left[\left(\frac{\partial \mathbf{z}}{\partial t}, \mathbf{w} \right) + \mathbf{a}(\boldsymbol{\xi}; \mathbf{z}, \mathbf{w}) \right] dt = \mathbf{0}, \quad \forall \mathbf{w} \in \mathbf{V} \times (0, T]. \quad (2.2)$$

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 [7]. In this paper we take \mathcal{T}_h to be a conforming finite element mesh consisting of simplices or parallelepipeds. Let the space of continuous, piecewise functions of degree q over the spatial domain be defined by

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

The semi-discrete variational formulation seeks, for any $\boldsymbol{\xi} \in \Gamma_{\boldsymbol{\xi}}$, $\mathbf{z}_h \in \mathbf{V}_h^{(q)} \times (0, T]$ such that,

$$\int_0^T \left[\left(\frac{\partial \mathbf{z}_h}{\partial t}, \mathbf{w} \right) + \mathbf{a}(\boldsymbol{\xi}; \mathbf{z}_h, \mathbf{w}) \right] dt = \mathbf{0}, \quad \forall \mathbf{w} \in \mathbf{V}_h^{(q)} \times (0, T]. \quad (2.3)$$

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 [9, 8].

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. 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, 1, 9]. 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:

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

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, 1, 9, 10] 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. 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. [13, 8, 1, 20], for a complete discussion of these additional terms.

4. Sparse grid approximations

Sparse grid stochastic collocation has been shown to provide efficient and accurate approximation of stochastic quantities [17, 18, 19]. 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 [16].

Sparse grids [3] approximate $J_h(\boldsymbol{\xi})$ via a weighted linear combination of basis functions

$$I_n[J_h(\boldsymbol{\xi})] := J_{h,n}(\boldsymbol{\xi}) = \sum_{k=1}^n v_k \Psi_k(\boldsymbol{\xi}) \quad (4.1)$$

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

$$\Psi_{\mathbf{l},\mathbf{i}}(\boldsymbol{\xi}) = \prod_{n=1}^d \psi_{l_n,i_n}(\xi_n) \quad (4.2)$$

where \mathbf{i} determines the location of a given grid point. There is a one-to-one relationship between Ψ_k in (4.1) 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 [3].

The multi-dimensional basis (4.2) span the discrete spaces $V_{\mathbf{1}} \subset L^2(\Gamma_{\boldsymbol{\xi}})$

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

These discrete spaces can be further decomposed into hierarchical difference spaces

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

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

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

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 $W_{\mathbf{1}}$ with an index set that satisfy

$$\mathcal{L} = \{ \mathbf{l} : |\mathbf{l}|_1 \leq l \} \tag{4.3}$$

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

$$J_{h,n}(\boldsymbol{\xi}) = \sum_{\mathbf{l} \in \mathcal{L}} J_{\mathbf{l}}, \quad J_{\mathbf{l}} = \sum_{\mathbf{i} \in \mathcal{I}_{\mathbf{l}}} v_{\mathbf{l},\mathbf{i}} \Psi_{\mathbf{l},\mathbf{i}}(\boldsymbol{\xi}) \tag{4.4}$$

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

Here we note that the $v_{\mathbf{l},\mathbf{j}}$ 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

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

Sparse grids can also be used to interpolate the forward solution \mathbf{z}_h and the adjoint solution ϕ_h . We define these approximations as $I_n[\mathbf{z}_h] := \mathbf{z}_{h,n}$ and $I_n[\phi_h] := \phi_{h,n}$. When computing a-posteriori error estimates,

interpolants of \mathbf{z}_h and ϕ_h must be computed at each point in the finite element mesh and step in the time discretization.

Although a vast improvement of full tensor grid approximations the number of sparse grid points in an isotropic sparse grid 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 [11], 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 considered important [14]. When using localized basis functions adaptivity can be performed combining the strengths of both dimension and regional adaptivity [17].

4.1. 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

$$\mathbf{l} - \mathbf{e}_k \in \mathcal{L} \text{ for } 1 \leq k \leq d, l_k > 1 \quad (4.5)$$

Each subspace that satisfies (4.5) 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. The auxiliary TERMINATE routine is given in Algorithm 2.

Purposefully, Algorithm 1 does not specify the contents of the INDICATOR and GLOBAL_INDICATOR routines in Algorithm 1. These routines control which subspaces are added to the sparse grid. The key to adaptive sparse grids working well is versions of 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 these indicator functions are functions of the hierarchical surplus values $v_{\mathbf{i},\mathbf{j}}$ of points in the grid. Throughout this paper we will use the indicator

$$\gamma_{\mathbf{l}} = \sum_{\mathbf{i} \in \mathcal{I}_{\mathbf{l}}} |v_{\mathbf{i},\mathbf{j}}| w_{\mathbf{i},\mathbf{j}} \quad (4.6)$$

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

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

It is possible alternative definitions of $\gamma_{\mathbf{l}}$ and η could produce more accurate sparse grids for a given number of function evaluations. Alternatives could be generated via studies such as that discussed in [23]. Such an exploration of alternative surplus based criteria, however, is beyond the scope of this paper. Instead this paper focuses on alternative criteria that do not require additional evaluations of the simulation model. These alternative criteria are discussed in Section 6. The differences between surplus refinement and the alternative refinement criteria proposed in Section 6 motivate the conditional statements in Algorithm 1 not usually seen in the literature [11, 17].

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

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

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

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

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

4.2. Local adaptivity

Locally-adapted sparse grids can be adapted using a 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.

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

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

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 [18] and the simultaneous dimension and local refinement proposed in [17] 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 $\xi_{\mathbf{l},i}$ in $\mathcal{A}_{\text{local}}$ with the largest error indicator $\gamma_{\mathbf{l},i}$. 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 [17].

Throughout this paper we will use the indicator

$$\gamma_{\mathbf{l},i} = |v_{\mathbf{i},j}| w_{\mathbf{i},j} \quad (4.7)$$

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

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

5. 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 which usually requires sampling the surrogate according to the distribution of the random parameters over Γ_{ξ} . 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 [5] and further analyzed in [4, 6] 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

$$J(\boldsymbol{\xi}) - J_{h,n}(\boldsymbol{\xi}) \approx \varepsilon(J_{h,n}(\boldsymbol{\xi})), \quad (5.1)$$

where

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

To construct the approximations of the forward and adjoint solutions we construct a sparse grid approximation for each degree of freedom in the entire forward and adjoint solutions. Note that unlike the sparse grid approximation of a scalar function given in (4.4), the approximations of $\mathbf{z}_{h,n}$ and $\phi_{h,n}$ are functions, not just of $\boldsymbol{\xi}$, but also the spatial and temporal mesh locations \mathbf{x} and \mathbf{t} . However, building approximations over \mathbf{x} and \mathbf{t} requires no additional PDE solves as one solve will produce an entire spatial and temporal history. We simply need to store and compute the hierarchical surplus for each \mathbf{x} and \mathbf{t} separately. This is not as memory intensive as it may first appear. We only need to store the additional hierarchical surpluses at each sparse grid point. The remaining sparse grid data structures do not need to be replicated. We do remark, however, that if the PDE solution has millions or even billions of degrees of freedom then memory problems associated with processing big data sets will arise. This issue is not just endemic to the approach proposed in this paper, but is a difficulty also faced by a posteriori error analysis for deterministic PDEs. The typical mitigation strategy to avoid saving the full forward approximation is called check-pointing [15, 22], which effectively balances the cost in saving the solution versus recomputing the solution. Unfortunately, this is not an option for the evaluation of the error estimate (5.1) since neither the forward problem nor the adjoint problem are actually solved at the sample point. We could utilize data compression to reduce the storage costs but this is beyond the scope of this paper and will be pursued in later work.

While the a posteriori error estimate given by (5.1) is only an approximation of the true error, it was shown in [4, 6] that the *error in the error estimate* can be bounded by the product of the pointwise errors (in $\Gamma_{\boldsymbol{\xi}}$) in the forward and adjoint approximation. Thus, for smooth problems the error estimate is higher order than the surrogate approximation and can be used as either a higher order correction term (as in [4, 6]) or to guide stochastic adaptivity as we pursue in Section 6.

6. Alternative adaptivity strategies

Hierarchical surplus indicators such as those in (4.6) and (4.7) require the model to be evaluated at the associated grid point $\boldsymbol{\xi}_{\mathbf{l},i}$ 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 (5.1) is one such indicator.

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

$$\gamma_{\mathbf{l}}^{\varepsilon} = \sum_{\mathbf{i} \in \mathcal{I}_{\mathbf{l}}} |\varepsilon(J_{h,n}(\boldsymbol{\xi}_{\mathbf{l},i}))| w_{\mathbf{i},j}, \quad \eta^{\varepsilon} = \sum_{\mathbf{l} \in \mathcal{A}} \gamma_{\mathbf{l}}^{\varepsilon} \quad (6.1)$$

and for locally adaptive sparse grids

$$\gamma_{\mathbf{l},i}^{\varepsilon} = |\varepsilon(J_{h,n}(\boldsymbol{\xi}_{\mathbf{l},i}))| w_{\mathbf{i},j}, \quad \eta^{\varepsilon} = \sum_{\{\mathbf{l},i\} \in \mathcal{A}} \gamma_{\mathbf{l},i}^{\varepsilon} \quad (6.2)$$

The use of these indicators requires a minor modification to the interpolate algorithms presented in the literature [11, 17]. 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 is evaluated at the sparse grid when `INDICATOR` is called on a subspace containing the point (dimension refinement) or on the point directly (local refinement). The point is then immediately added to the sparse grid. 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.

When using local refinement one should ensure all the ancestors of and point in the sparse grid also exist in the sparse grid. Similarly to when deciding to evaluate the model at a point, the use of a posteriori error guided refinement also effects the point in `INTERPOLATE` that the ancestors are added to the sparse grid. Although the function values at ancestor points are not necessary to compute the error estimate at the child point we add all ancestors to the sparse grid before the error estimate is computed. Alternatively one can also add the ancestor point later, only when the point is added to the grid. However 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 for all of the examples presented in this paper. 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 [18] and generalized local refinement [17] 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 generalized local refinement.

7. Sparse grid approximations of the error

As previously mentioned, the evaluation of the a posteriori error estimate using (5.1) requires the full forward and adjoint approximations at $\boldsymbol{\xi} \in \Gamma_{\boldsymbol{\xi}}$ as well as the evaluation of the space-time weighted-residual. While this is usually much cheaper than solving the forward problem at $\boldsymbol{\xi}$, the computation cost in producing the error estimate at $\boldsymbol{\xi}$ 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 (5.1).

Given a sparse grid approximation of the quantity of interest $J_{h,n}(\boldsymbol{\xi})$ we can use a finite number of error estimates $\varepsilon(J_{h,n}(\boldsymbol{\xi}_{1,i}))$ 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 (5.1).

The procedure needed to construct an enhanced sparse grid is outlined in Algorithm 3. To build an enhanced sparse grid we must first construct sparse grid approximations of the QOI, forward solution and the adjoint solution. The number of points used in these approximations is controlled by two parameters τ and n . The tolerance τ should be chosen such that refinement is stopped when the interpolation error in the enhanced sparse grid approximation J_n^ε becomes less than the *error in the error estimate*. Setting the maximum number of points n to a finite number can also be used as a means to limit the cost of constructing the enhanced sparse grid

to within a pre-defined computational budget. The choices of τ and n found in Algorithm 3 are discussed in Section 7.2. Once the approximations of the QOI, forward solution and the adjoint solution have been obtained the sparse grid is enhanced by sampling the error estimate $J_n^\varepsilon(\boldsymbol{\xi})$. These two phases needed to build the enhanced approximation $J_{n,m}^\varepsilon$ are graphically illustrated in Figure 1.

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

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

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

% $\mathbf{z}_{h,n/4}$ and $\phi_{h,n/4}$ are needed to compute $\varepsilon(J_{h,n}(\boldsymbol{\xi}))$.

An enhanced sparse grid can be built using dimension or local refinement. The construction of the enhanced sparse grid approximation varies slightly depending on whether local or dimension adaptivity is being used. The two types of construction are discussed in Section 7.1.

7.1. Adaptivity

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

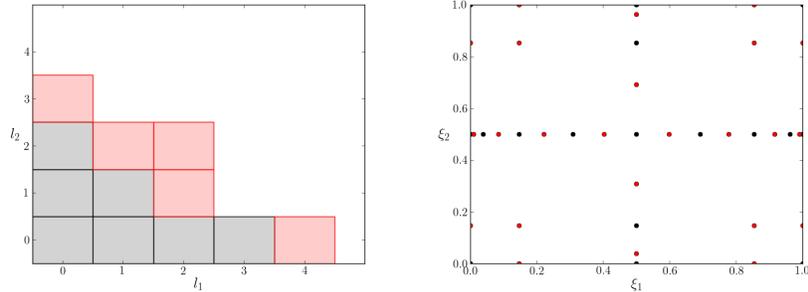


Figure 1: The two phases undertaken when building an enhanced dimension adaptive sparse grid interpolant J_n^ε . The first phase consist of building the un-enhanced approximation $J_{h,n}(\boldsymbol{\xi})$. The gray boxes on the left represent the 2D sparse grid subspaces $\mathbf{l} = (l_1, l_2)$ used to build $J_{h,n}(\boldsymbol{\xi})$ 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 $J_{n,m}^\varepsilon$. 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 $J_{h,n}(\boldsymbol{\xi})$. The red boxes are not to be confused with active indices which cannot have the structure shown in this figure.

Given the sets \mathcal{A} and \mathcal{L} we simply use $\text{INTERPOLATE}[J_n^\varepsilon(\boldsymbol{\xi}), \mathcal{L}, \mathcal{A}, \tau, m]$ to construct the enhanced sparse grid $J_{n,m}^\varepsilon := I_m[J_n^\varepsilon(\boldsymbol{\xi})]$, where $J_n^\varepsilon(\boldsymbol{\xi}) := J_{h,n}(\boldsymbol{\xi}) + \varepsilon(J_{h,n}(\boldsymbol{\xi}))$. Unlike the construction of $J_{h,n}(\boldsymbol{\xi})$ the construction of J_n^ε does not require the evaluation of the forward solution $\mathbf{z}_h(\boldsymbol{\xi}_{\mathbf{l},i})$ at each new grid point. Rather, evaluating $J_n^\varepsilon(\boldsymbol{\xi})$ only involves evaluating the un-enhanced sparse grid $J_{h,n}(\boldsymbol{\xi}_{\mathbf{l},i})$ and computing the error estimate $\varepsilon(J_{h,n}(\boldsymbol{\xi}_{\mathbf{l},i}))$. 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 $\text{INTERPOLATE}[J_n^\varepsilon(\boldsymbol{\xi}), \mathcal{L}, \mathcal{A}, \tau, m]$ is called the function value at all points in \mathcal{L} must be updated to include the physical discretization error estimate $\delta(J_h(\boldsymbol{\xi}))$. That is for all $\boldsymbol{\xi}_{\mathbf{l},i} \in \mathcal{L}$ set the function value at that point to $J_h(\boldsymbol{\xi}_{\mathbf{l},i}) + \delta(J_h(\boldsymbol{\xi}_{\mathbf{l},i}))$. The hierarchical surpluses at each of these points must also be update accordingly. We do this because any new points to the enhanced sparse grid will have function values $J_{h,n}(\boldsymbol{\xi}_{\mathbf{l},i}) + \varepsilon(J_{h,n}(\boldsymbol{\xi}_{\mathbf{l},i}))$, where $\varepsilon(J_{h,n}(\boldsymbol{\xi}_{\mathbf{l},i}))$ 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 constructing 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} and \mathcal{A} and then call INTERPOLATE after adjusting the function values and hierarchical surpluses for all $\xi_{1,i} \in \mathcal{L}$.

7.2. *Balancing the stochastic and physical discretization errors*

When quantifying uncertainty of models, the deterministic and stochastic discretization error must both 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 physical discretization error is pre-determined and cannot be changed. When balancing stochastic error with a fixed physical discretization error with this assumption, we must be able to handle two regimes: firstly that our computational budget is large enough to drive the stochastic error below the physical discretization error; and secondly it will not be possible to 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. So when $J_n^\varepsilon(\xi)$ is constructed, we set τ to be the maximum of the approximate physical discretization error δ_{\max} and the approximate stochastic error γ_{\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 [5, 4, 6]. 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 η^2 .

In the second regime, the goal is to reduce the total error as much as possible within the computational budget. To do this we must consider the costs of solving the forward problem, solving the adjoint problem, and evaluating the error estimate (5.1). 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 $\varepsilon(J_{h,n}(\xi))$ to build the enhanced sparse grid. In many cases, using half the computational budget for computing $\varepsilon(J_{h,n}(\xi))$ is overkill. This is especially true if the cost in evaluating the error estimate is 2-3 orders of magnitude cheaper than solving the forward or adjoint problems. Consequently, we use an error indicator to terminate the sparse grid earlier if a desired tolerance

is met. Note 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.

8. Results

In this section, we provide several numerical examples to illustrate the properties and convergence of the proposed methodology. All of the examples use 100,000 Latin Hypercube samples to compute the discrete ℓ_2 error, unless otherwise stated. In all of the figures, 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 §6.

8.1. Parameterized linear system

Consider the diagonally dominant parameterized linear system

$$\begin{bmatrix} \xi_1 & \xi_2 \\ \xi_3 & \xi_4 \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \end{bmatrix} \quad \text{or in matrix notation} \quad \mathbf{A}\mathbf{z} = \mathbf{b} \quad (8.1)$$

with $\xi_1, \xi_3 \in [2, 4]$ and $\xi_2, \xi_4 \in [0, 1]$. The adjoint equation is given by

$$\mathbf{A}^T \phi = \psi \quad (8.2)$$

We take our quantity of interest to be the first component of the solution \mathbf{z} , so we set $\psi = [1, 0]^T$.

As an initial step to illustrating the benefits of the proposed methodology, we mirror the work outlined in [4] and demonstrate the accuracy gained by enhancing the interpolant of the linear functional $J_{h,n}(\boldsymbol{\xi})$ with samples from the error estimate $\varepsilon(J_{h,n}(\boldsymbol{\xi}))$.

We build a surplus-based dimension-adaptive sparse grid using Clenshaw-Curtis abscissa and compute $\|J_h(\boldsymbol{\xi}) - J_{h,n}(\boldsymbol{\xi})\|_{\ell_2(\Gamma_{\boldsymbol{\xi}})}$ by sampling $\boldsymbol{\xi}$ using a 100,000 point Latin hypercube. We then compute error-estimates at the same 100,000 points and calculate $\|J_h(\boldsymbol{\xi}) - (J_{h,n}(\boldsymbol{\xi}) + \varepsilon(J_{h,n}(\boldsymbol{\xi})))\|_{\ell_2(\Gamma_{\boldsymbol{\xi}})}$. Note that this example does not contain any physical discretization error, so $J(\boldsymbol{\xi}) = J_h(\boldsymbol{\xi})$.

Figure 2 shows the ℓ^2 error in the functional value and the error in the improved function value as the computational cost increases. 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.

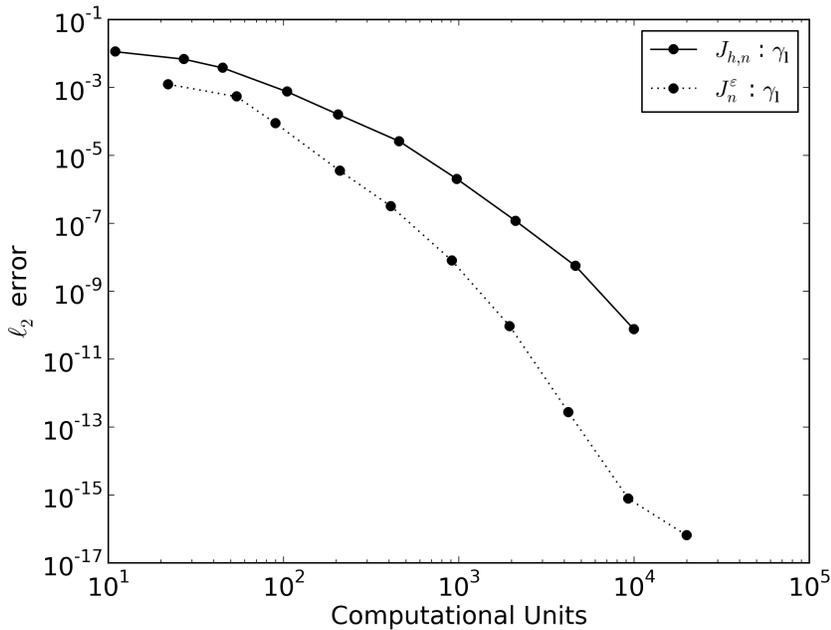


Figure 2: Approximation of the 4d parameterized linear system using a dimension-adaptive Clenshaw-Curtis sparse grid.

Figure 3 compares error in the functional value and the error in the improved function value as the computational cost increases using local adaptivity based upon one-dimensional equidistant meshes and piecewise quadratic basis functions using surplus-based refinement. Throughout the remainder of this manuscript, a quadratic basis will be used when local adaptivity is employed, unless otherwise stated.

Similarly to the results obtained in [4], which were obtained using enhanced polynomial chaos expansions, enhancing both dimension adaptive and locally adaptive sparse grids with error estimates greatly increases accuracy for a fixed computational cost.

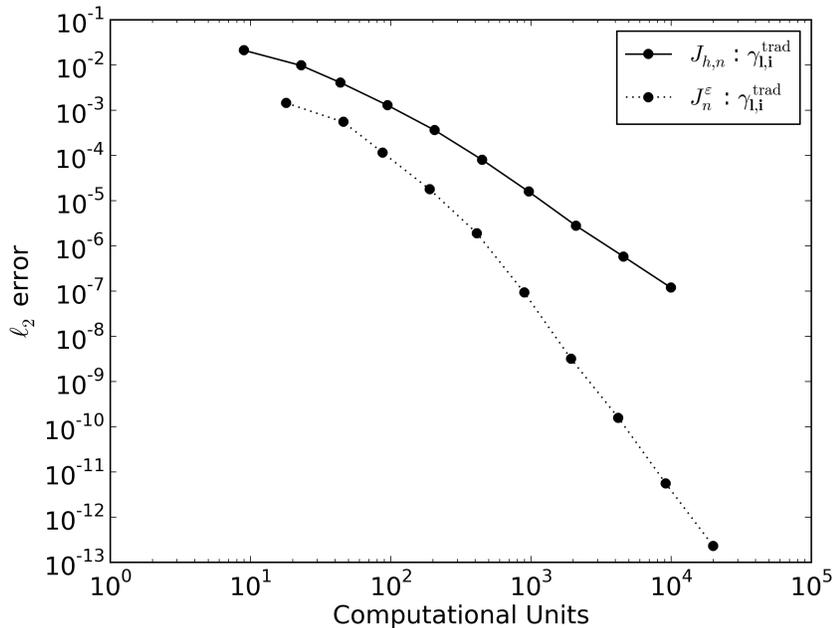


Figure 3: Approximation of the 4d parameterized linear system using a locally-adaptive sparse grid with a quadratic basis

8.2. 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:

$$-\frac{d}{dx} \left[a(x, \boldsymbol{\xi}) \frac{d\mathbf{z}}{dx}(x, \boldsymbol{\xi}) \right] = 10, \quad (x, \boldsymbol{\xi}) \in (0, 1) \times I_{\boldsymbol{\xi}} \quad (8.3)$$

subject to the physical boundary conditions

$$\mathbf{z}(0, \boldsymbol{\xi}) = 0, \quad \mathbf{z}(1, \boldsymbol{\xi}) = 0 \quad (8.4)$$

Furthermore assume that the log random diffusivity satisfies

$$a(x, \boldsymbol{\xi}) = \bar{a} + \sigma_a \sum_{k=1}^d \sqrt{\lambda_k} \phi_k(x) \xi_k \quad (8.5)$$

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

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

The variability of the diffusivity field (8.5) is controlled by σ_a and the correlation length l_c which determines the decay of the eigenvalues λ_k . Here we set $d = 25$, $\sigma_a = 1.$, $l_c = 0.1$, $\bar{a} = 0$ and $\xi_k \in [-1, 1]$, $k = 1, \dots, d$ to be independent and uniformly distributed random variables. We define our quantity of interest to be the value of the solution at $x=0.5$ which we approximate using the linear functional $J(\boldsymbol{\xi})(\boldsymbol{\xi}) = (\psi(x), \mathbf{z}(x, \boldsymbol{\xi}))$ where

$$\psi(x) = C_s \exp(-100(x - .5)^2),$$

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

$$-\frac{d}{dx} \left[a(x, \boldsymbol{\xi}) \frac{d\phi}{dx}(x, \boldsymbol{\xi}) \right] = \psi(x), \quad (x, \boldsymbol{\xi}) \in (0, 1) \times I_{\boldsymbol{\xi}} \quad (8.6)$$

subject to the physical boundary conditions

$$\phi(0, \boldsymbol{\xi}) = 0, \quad \phi(1, \boldsymbol{\xi}) = 0 \quad (8.7)$$

Unlike the parameterized linear system discussed in the previous section, the solution to (8.3) must be approximated numerically which introduces significant physical discretization error (error labeled I in (1.1)). This physical discretization error must be accounted for when quantifying uncertainty. We approximate the solution to (8.3) using a standard continuous Galerkin finite element method on a uniform mesh with 100 elements ($h = 0.01$) and piecewise linear basis functions. To include the deterministic component in the error estimate, we need to solve the adjoint problem using a higher-order discretization than was used to solve the forward problem. In this paper, we use a continuous Galerkin method with piecewise quadratic basis functions to approximate the adjoint solution. The deterministic error using the aforementioned discretization is approximately $\delta_{\max} = 8.41 \times 10^{-5}$. Any computational effort used to drive the stochastic error below the deterministic error will be wasted. Figure 4 compares the convergence of an un-enhanced sparse surplus-based, dimension-adaptive, Clenshaw-Curtis grid $J_{h,n}(\boldsymbol{\xi})$ with the enhanced sparse grid approximation J_n^{ε} . Again we have computed the error estimates directly at the 100,000 point Latin hypercube samples used to

compute $\|J_h(\boldsymbol{\xi}) - J_{h,n}(\boldsymbol{\xi})\|_{\ell_2(\Gamma_\xi)}$ and $\|J_h(\boldsymbol{\xi}) - (J_{h,n}(\boldsymbol{\xi}) + \varepsilon(J_{h,n}(\boldsymbol{\xi})))\|_{\ell_2(\Gamma_\xi)}$. Clearly, again there is significant benefit from using error estimates to improve the sparse grid interpolant.

Figure 4 also shows the importance of not refining the sparse grid when the stochastic error is smaller than the physical discretization error. 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. Note this stagnation does not occur in Figures 2 and 3 because the parameterized linear system being studied does not have a physical discretization error.

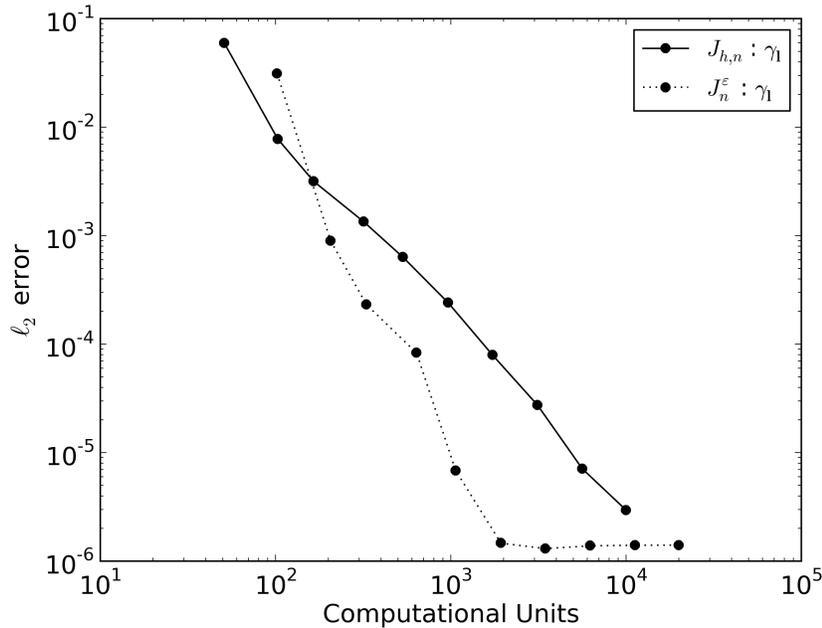


Figure 4: Approximation of the 25d heat equation using a dimension-adaptive Clenshaw-Curtis sparse grid

The cost of computing $\varepsilon(J_{h,n}(\boldsymbol{\xi}))$ is non-trivial and in this example the cost of evaluating $\varepsilon(J_{h,n}(\boldsymbol{\xi}))$ at 100,000 points would outweigh the cost of the forward and adjoint solves. Consequently we must build an approximation of the error estimate using the approach described in Section 7. In all plots 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 ap-

proximation, and calculating the refinement indicators. In all of these plots, we also include J_n^ε for reference, but this data does not include the computational cost in computing the error estimate. In this example, we assume that the cost of evaluating the error estimate is 1/25 of a computational unit.

The need to set appropriate termination criteria outlined in Section 7.2 is illustrated in Figure 5. The figure illustrates that as more forward and adjoint solutions are used to construct $J_{h,n}(\xi)$ the higher accuracy that can be obtained in the enhanced sparse grid $J_{n,m}^\varepsilon$. However the accuracy that can be obtain is limited, yet any error indicator that does not know the ‘true’ values of the quantity of interest will appear to converge to machine precision. Consequently, as stated in Section 7.2, we set τ when constructing the enhanced sparse grid in Algorithm 3 to be maximum of the approximate physical discretization error δ_{\max} and the approximate stochastic error γ_{\max} in the enhanced sparse grid.

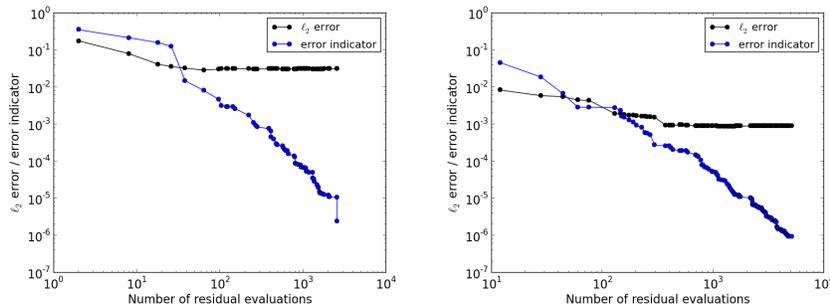


Figure 5: The decay of the ℓ_2 error vs the decay in the error indicator for the enhanced sparse grid $J_{n,m}^\varepsilon$, starting from an initial grid with 51 points (left) and 103 points (right).

Figure 6 plots the convergence of the un-enhanced surplus-based dimension-adaptive Clenshaw-Curtis sparse grids $J_{h,n}(\xi)$, J_n^ε and $J_{n,m}^\varepsilon$. Both the un-enhanced sparse grid and the enhanced sparse grid are refined until the estimate of the stochastic error γ_{\max} is below the deterministic error δ_{\max} . In the case of the un-enhanced grid this point is reached, and the convergence study is terminated, when the sparse grid has 10,000 points. The computational cost of $J_{n,m}^\varepsilon$ includes the cost of evaluating the residual to compute the error estimates needed at the sparse grid points. Here, we assume that the cost of an error estimate is 1/25 of the cost of one forward solve.

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 . This is not a fair comparison as in practice, using J_n^ε is typically computationally infeasible. However the comparison does show that the choice of termination conditions in Algorithm 1 means that the grid terminates before the error in the enhanced grid is below the error in the error estimate. As Figure 5 shows this level of accuracy can be obtained however it comes at the cost of (possibly many) additional residual computations.

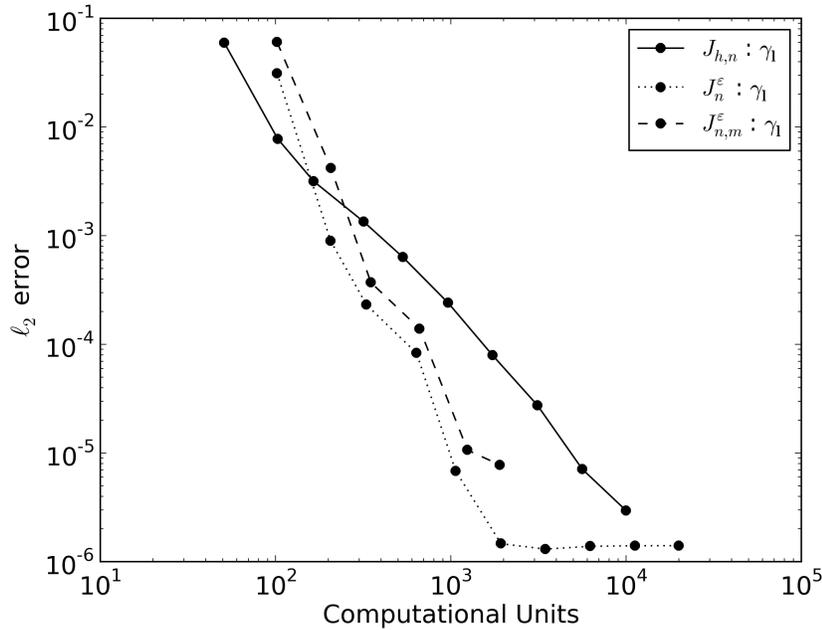


Figure 6: Comparison of the ℓ_2 convergence of the enhanced dimension-adaptive sparse grid using direct evaluation of the error estimate J_n^ε and a sparse grid approximation of the enhanced QOI $J_{n,m}^\varepsilon$.

Figure 7 plots the convergence of the un-enhanced surplus-based locally-adaptive sparse grids $J_{h,n}(\xi)$, 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 $J_{h,n}(\xi)$ is reduced however this can be addressed by the use of different refinement strategies discussed in Section 8.4

8.3. Non-linear coupled system of ODEs

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

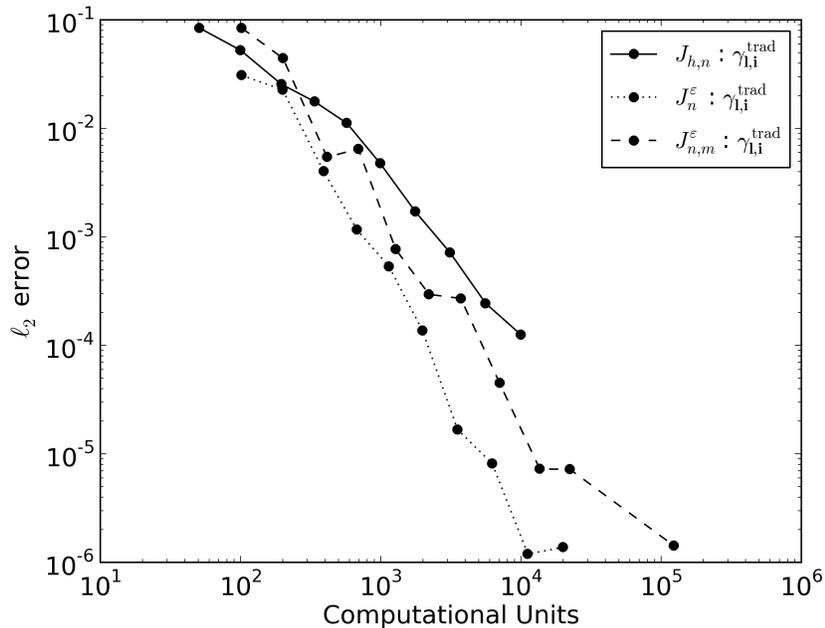


Figure 7: Comparison of the ℓ_2 convergence of the enhanced locally-adaptive sparse grid using direct evaluation of the error estimate and a sparse grid approximation of the enhanced QOI

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

$$\begin{cases} \frac{dz_i}{dt} = r_i z_i \left(1 - \sum_{j=1}^3 \alpha_{ij} z_j \right), & t \in (0, 10], \\ \mathbf{z}_i(0) = \mathbf{z}_{i,0} \end{cases}, \quad (8.8)$$

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 (8.8) 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

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

$$\begin{cases} -\frac{d\phi_i}{dt} = r_i\phi_i \left(1 - \sum_{j=1}^3 \alpha_{ij}\mathbf{z}_j\right) + r_i\mathbf{z}_i \left(1 - \sum_{j=1}^3 \alpha_{ji}\phi_j\right), & t \in (10, 0], \\ \phi_i(10) = 0, & i = 1, 2 \\ \phi_3(10) = 1. \end{cases}, \quad (8.9)$$

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

Figure 8 and Figure 9 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 $J_{h,n}(\boldsymbol{\xi})$. As seen in the previous example, the accuracy of $J_{n,m}^\varepsilon$ is not as high as J_n^ε but could be improved at the cost of additional residual calculations (error estimates) at new sparse grid points.

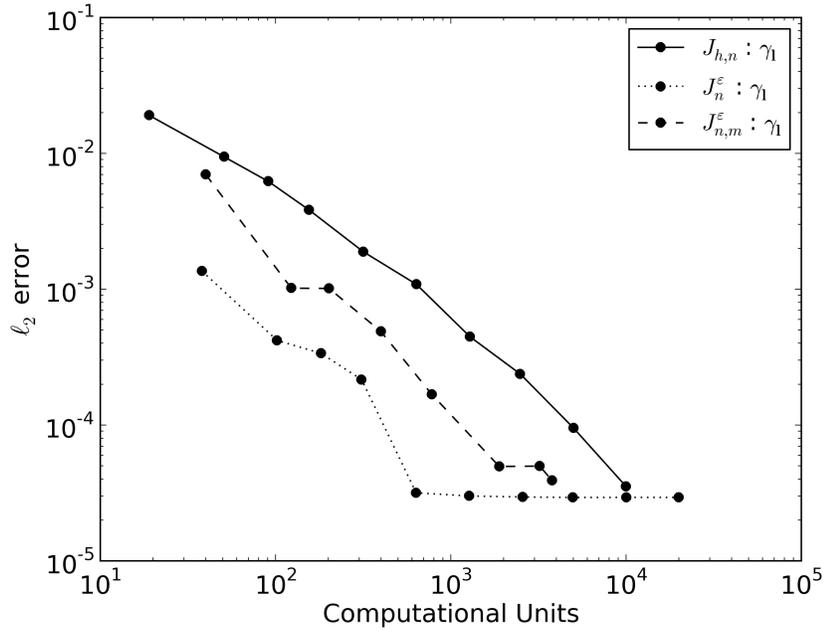


Figure 8: Comparison of the ℓ_2 convergence of the enhanced dimension-adaptive sparse grid using direct evaluation of the error estimate and a sparse grid approximation of the enhanced QOI

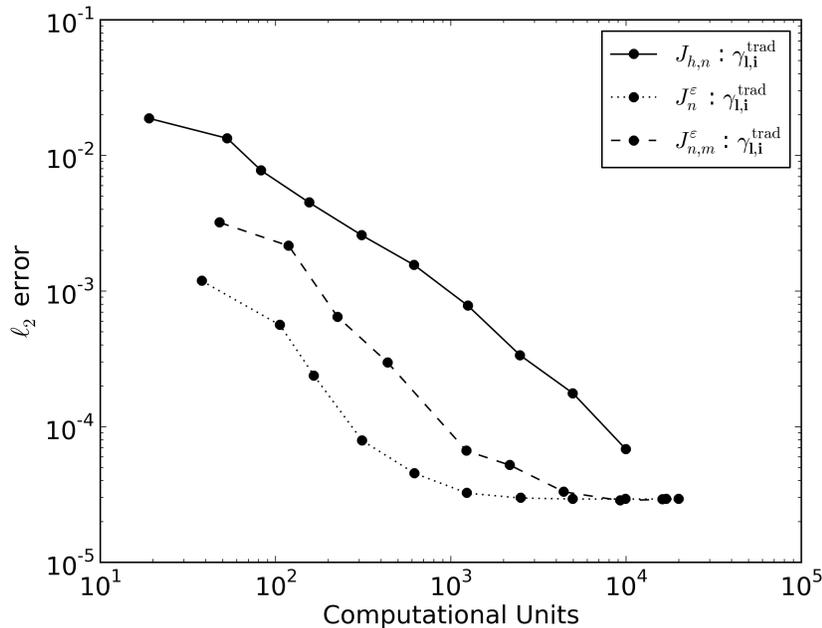


Figure 9: Comparison of the ℓ_2 convergence of the enhanced locally-adaptive sparse grid using direct evaluation of the error estimate and a sparse grid approximation of the enhanced QOI

8.4. Alternative Refinement Strategies

In this section we demonstrate the utility of using the alternative refinement strategies outlined in Section 6. 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.

Figure 10 compares refinement strategies for the 25d heat equation problem presented in Section 8.2. 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. Thus the function evaluations that are needed by surplus refinement that 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 reduce. This point is illustrated in Figure 11 which compares a posteriori and surplus dimension adaptive sparse grids when used to solve the 9d predator-prey model presented in Section 8.3.

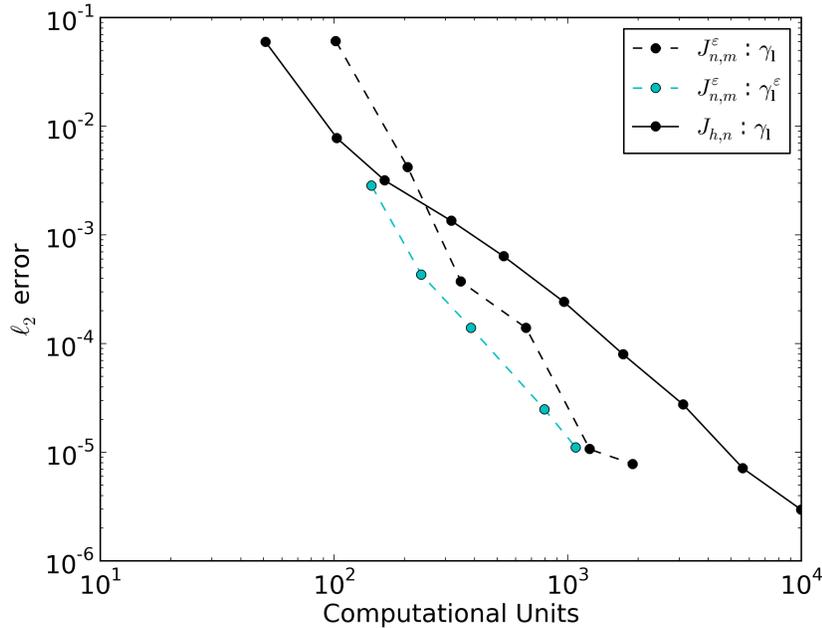


Figure 10: Comparison of dimension-adaptive refinement strategies when applied to the 25d heat equation

A posteriori refinement can also be used in conjunction with local refinement. Figure 12 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 13 compares refinement strategies for the 9d predator prey model. Due to the low degree of anisotropy of this function generalized local refinement provides now improvement over traditional local refinement. In the case of the surplus-based enhanced sparse grid the accuracy is slightly worse.

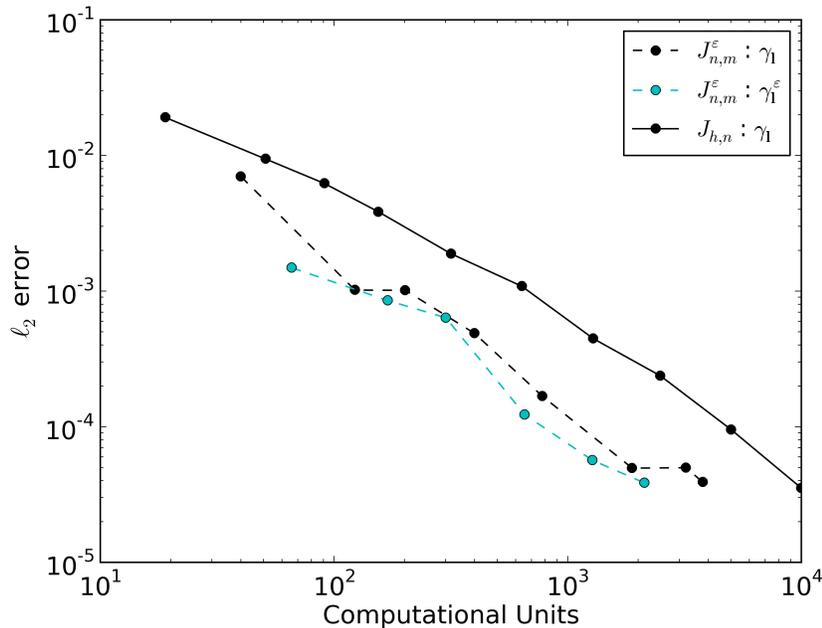


Figure 11: Comparison of dimension-adaptive refinement strategies when applied to 9d predator prey model

Traditional surplus based refinement is the most accurate strategy here, but the a posteriori error based generalized local refinement has comparable accuracy.

9. 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 and to drive adaptivity. 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.

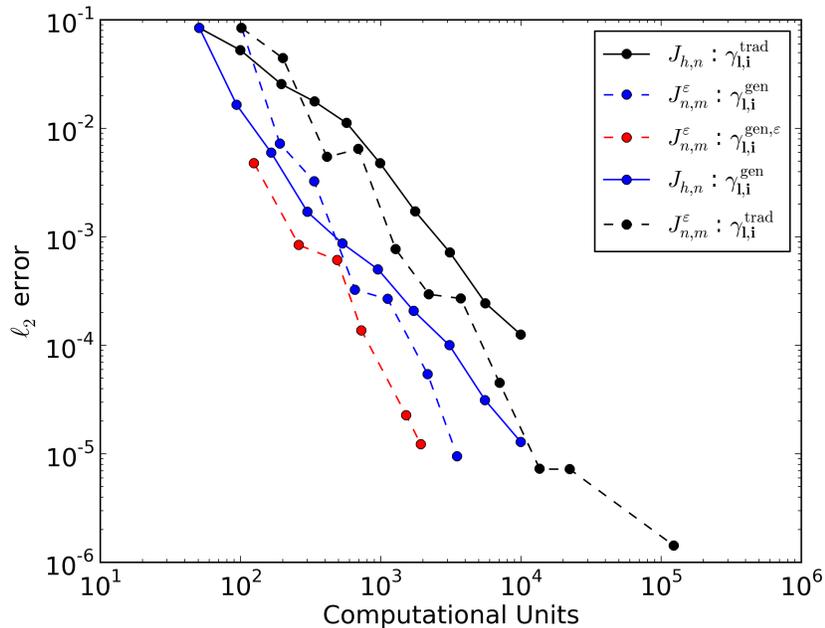


Figure 12: Comparison of local refinement strategies when applied to the 25d heat equation

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 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.

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.

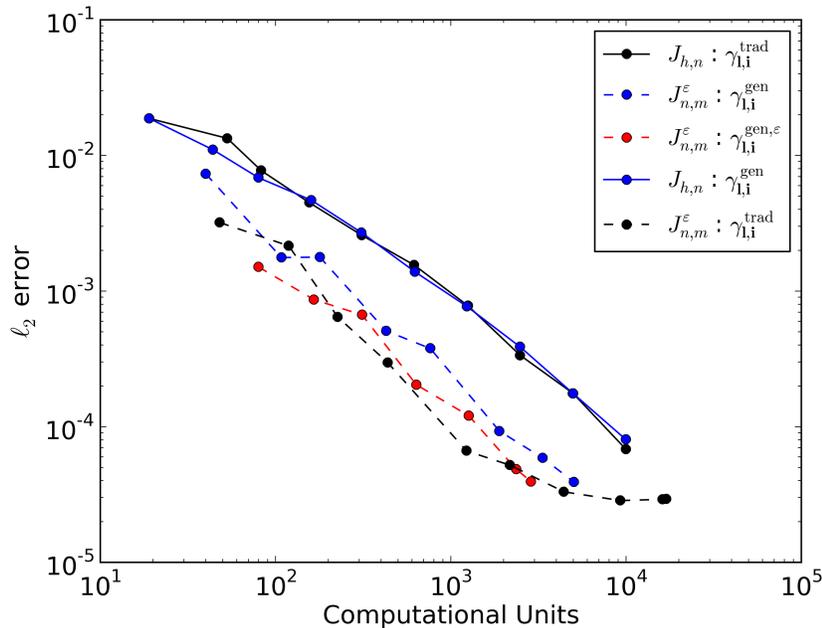


Figure 13: Comparison of local refinement strategies when applied to the 9d predator prey model

References

- [1] BANGERTH, W., AND RANNACHER, R. *Adaptive Finite Element Methods for Differential Equations*. Birkhauser Verlag, 2003.
- [2] BECKER, R., AND RANNACHER, R. An optimal control approach to a posteriori error estimation in finite element methods. *Acta Numerica* 10 (2001), 1–102.
- [3] BUNGARTZ, H.-J., AND GRIEBEL, M. Sparse grids. *Acta Numerica* 13 (2004), 147–269.
- [4] 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.
- [5] BUTLER, T., DAWSON, C., AND WILDEY, T. A posteriori error analysis of stochastic spectral methods. *SIAM J. Sci. Comput.* 33 (2011), 1267–1291.

- [6] 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.
- [7] CIARLET, P. G. Basic error estimates for elliptic problems. In *Handbook of numerical analysis, Vol. II*. North-Holland, Amsterdam, 1991, pp. 17–351.
- [8] ERIKSSON, K., ESTEP, D., HANSBO, P., AND JOHNSON, C. *Computational differential equations*. Cambridge University Press, Cambridge, 1996.
- [9] 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.
- [10] 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.
- [11] GERSTNER, T., AND GRIEBEL, M. Dimension-adaptive tensor-product quadrature. *Computing* 71, 1 (SEP 2003), 65–87.
- [12] GHANEM, R., AND SPANOS, P. *Stochastic Finite Elements: A Spectral Approach*. Springer Verlag, New York, 2002.
- [13] 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.
- [14] GRIEBEL, M. Adaptive sparse grid multilevel methods for elliptic PDEs based on finite differences. *Computing* 61, 2 (1998), 151–179.
- [15] GRIEWANK, A., AND WALTHER, A. Algorithm 799: revolve: an implementation of checkpointing for the reverse or adjoint mode of computational differentiation. *ACM Trans. Math. Softw.* 26, 1 (Mar. 2000), 19–45.
- [16] JAKEMAN, J., ELDRÉD, M., AND XIU, D. Numerical approach for quantification of epistemic uncertainty. *J. Comput. Phys.* 229, 12 (2010), 4648–4663.

- [17] 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.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] RASMUSSEN, C. E. Gaussian processes for machine learning.
- [22] WANG, Q., MOIN, P., AND IACCARINO, G. Minimal repetition dynamic checkpointing algorithm for unsteady adjoint calculation. *SIAM J. Sci. Comput.* *31*, 4 (2009), 2549–2567.
- [23] 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.
- [24] XIU, D., AND KARNIADAKIS, G. The Wiener-Askey Polynomial Chaos for stochastic differential equations. *SIAM J. Sci. Comput.* *24*, 2 (2002), 619–644.