An adaptive method for solving stochastic equations based on interpolants over Voronoi cells

Wayne Isaac T. Uy
Center for Applied Mathematics,
Cornell University
wtu4@cornell.edu

Mircea D. Grigoriu
School of Civil and Environmental Engineering and
Center for Applied Mathematics,
Cornell University
mdg12@cornell.edu

Abstract

An adaptive collocation-based surrogate model is developed for the solution of equations with random coefficients, referred to as stochastic equations. The surrogate model is defined on a Voronoi tessellation of the samples of the random parameters with centers chosen to be statistically representative of these samples. We investigate various interpolants over Voronoi cells in order to formulate surrogates and analyze their convergence properties. Unlike Monte Carlo solutions, relatively small numbers of deterministic calculations are needed to implement surrogate models. These models can be used to generate large sets of solution samples with a minimum computational effort. In this work, we propose a framework for an adaptive construction of the surrogate such that by refining the Voronoi cells, the mapping between the random parameters and the solution is incorporated. A rigorous refinement measure which is quantitatively indicative of the performance of the surrogate is used to drive adaptivity. We present numerical examples that compare this surrogate with other collocation-based surrogates and demonstrate the theoretical aspects of the adaptive method.

1 Introduction

We consider the mapping $Z \mapsto U(x,t,Z)$ where Z represents a random vector defined on the probability space $(\Omega, \mathcal{F}, P)$ while x and t represent the spatial and temporal variables, respectively. Such mappings arise in applications, for instance problems in mechanics [8], in

which the response $U(x,t,Z)$ solves the forward stochastic equation $\mathcal{L}(U(x,t,Z)) = 0$ where the operator $\mathcal{L}$ typically characterizes a stochastic (partial) differential equation wherein Z encodes the uncertainty in the parameters of the physical model.

In order to quantify the behavior of the stochastic response $U(x,t,Z)$, probabilistic quantities of interest such as moments of U, distribution functions, probabilities of failure, etc., are sought after in applications. The only general method to accomplish this task is Monte Carlo. The method proceeds by generating a large number of samples of Z and solving the forward equation for each sample. As such, the implementation is straightforward and the computational cost does not scale with the dimension of the random vector Z. However, the method poses a computational burden for thousands of response calculations. To ameliorate the computational demand of Monte Carlo, surrogate models of the response $U(x,t,Z)$ have been introduced in which an approximation of the response is constructed in the image space of the random vector, $\Gamma = Z(\Omega) \in \mathbb{R}^d$. More specifically, a surrogate model is a mapping $Z \mapsto \widetilde{U}(x,t,Z)$ in which $\widetilde{U}(x,t,Z)$ approximates $U(x,t,Z)$ in some sense and for which samples of the response are cheaper to procure. We will drop the spatial and temporal variables for notational convenience in what follows and assume that $U(Z)$ is real-valued for simplicity.

Two general methods for constructing surrogate models for stochastic equations have flourished over the years, namely polynomial chaos or stochastic Galerkin method [14, 15] and stochastic collocation method [5–11, 19]. Stochastic Galerkin method projects the response $U(Z)$ onto a set of polynomial basis functions which are orthogonal with respect to a probability measure. While the resulting polynomial surrogate $\widetilde{U}(Z)$ is convenient to evaluate, $\widetilde{U}(Z)$ only converges to $U(Z)$ in $L^2(\Omega)$ and samples of $\widetilde{U}(Z)$ are prone to oscillatory behavior, cf. [13, 17]. The lack of guarantee for $L^p(\Omega)$ convergence for $p > 2$ implies that $\widetilde{U}(Z)$ is impractical for use in extreme events modeling. In addition, even though $L^2(\Omega)$ convergence guarantees convergence in distribution, the rate of convergence of the latter can be slow so that truncated levels, which needs to be kept low in applications for computational reasons, may yield unsatisfactory approximations. Stochastic collocation, on the other hand, proceeds by selecting a finite number of collocation or interpolation nodes $z_k \in \mathbb{R}^d$ in $\Gamma = Z(\Omega)$, evaluating $U(z_k)$, and constructing an interpolant $\widetilde{U}(z)$ over Γ. Smolyak-based sparse grid collocation has been proposed [9, 10] to overcome the curse of dimensionality when Z is high-dimensional. The performance of stochastic collocation hinges on the choice of interpolating functions and the location of the collocation nodes, aside from the regularity of the response $U(Z)$. It has been shown in [19] and in [10, Theorem 3.7] that under some assumptions on $U(Z)$, and for appropriate choices of the interpolant and collocation nodes, $\widetilde{U}(Z)$ converges to $U(Z)$ in $L^\infty(\Omega)$. A drawback of this method is that the quality of $\widetilde{U}(Z)$ may deteriorate as the number of collocation nodes increases [12] and that, except for [7, 18], the nodes are oftentimes selected without regard to the probability law of Z.

In an effort to improve the performance and accuracy of surrogate models, adaptive methods and domain decomposition techniques have been implemented which capture the behavior

of the response. An adaptive procedure for stochastic collocation was used in [19] wherein collocation nodes are sequentially incorporated into the surrogate. The hierarchical surplus, which is the difference between the response and the interpolant at the preceeding iteration, is used as a criterion to determine which set of nodes will be prioritized at each iteration. This surplus serves as an indicator of where $U(Z)$ exhibits substantial variation in Γ. Other variations in the indicator include the hierarchical surplus weighted by the probability density function value at a node as in [7] as well adjoint-based a posteriori error estimates in place of the surplus as in [5].

Domain decomposition techniques have also been applied in conjunction with adaptive methods to obtain local surrogate models in subdomains of the image of the random vector Z, denoted by Γ, thereby addressing the deficiencies of global surrogates. For example, the works [11, 14, 15] employed rectangular meshes to partition Γ after which polynomial chaos or stochastic collocation method is performed in each partition to obtain a local surrogate $\widetilde{U}_k(Z)$. Refinement criteria were then proposed to determine which partitions to refine. In these works, the decay rate of the variance of $\widetilde{U}_k(Z)$ was used as the refinement measure while partitions chosen to be refined are evenly split along the dimensions of Z which are deemed important. Performing such procedure in high dimensions, however, is challenging as refining a rectangular partition in this manner yields a large number of offspring partitions. Furthermore, rectangular partitions are not well-suited for partitioning non-hypercube domains Γ. In contrast to this type of domain decomposition, the authors in [3, 4] pursue a partitioning of Γ via a Delaunay triangulation based on Newton-Cotes nodes in Γ. The refinement criterion for a simplex partition is then based on a very crude approximation of the local $L^2(\Omega)$ error between $U(Z)$ and $\widetilde{U}_k(Z)$. A major difficulty of this approach in high dimensions is that 2^d+1 evaluations of the forward model are needed to interpolate over the initial triangulation, where d is the dimension of Z. We further remark that the refinement and stopping criteria considered in the aforementioned works do not clearly quantify the performance of the surrogate model being refined.

To address some of the difficulties associated with existing surrogate models, we employ a novel surrogate model $\widetilde{U}(Z)$ introduced in [2]. As in stochastic collocation, a finite number of interpolation nodes $z_k \in \Gamma$ is selected where the set $\{z_k\}$ constitutes a Stochastic Reduced Order Model (SROM) of the random vector Z. Unlike stochastic collocation, the set $\{z_k\}$ is chosen such that it captures statistical properties of Z. A Voronoi tessellation on the samples of Z is then constructed using $\{z_k\}$ as centers and this tessellation establishes a partitioning of the image $\Gamma \in \mathbb{R}^d$. On each Voronoi cell, a local surrogate model $\widetilde{U}_k(Z)$ is formulated by performing a Taylor expansion of $U(Z)$ at $Z = z_k$. As demonstrated in [2], $\widetilde{U}(Z) \to U(Z)$ in $L^p(\Omega)$ and almost surely under mild conditions, and that the computational cost scales linearly with the dimension of Z. In addition, the surrogate response $\widetilde{U}(Z)$ can be conveniently utilized to estimate probabilistic quantities of interest related to the response. We observe, however, that because the mapping $Z \mapsto U(Z)$ is not incorporated in partitioning Γ then two mappings $Z \mapsto V(Z)$ and $Z \mapsto W(Z)$ will yield surrogates with identical Voronoi tessellations of Γ for identical Z, thereby neglecting to account for the

variations unique to each mapping.

We henceforth propose a framework for an adaptive method for the SROM-based surrogate described in [2] which takes into account the mapping $Z \mapsto U(Z)$ in the domain decomposition. The method sequentially builds a SROM $\{\widetilde{z}_k\}$ for Z that prioritizes regions of Γ with high probability and regions which manifest sharp variations and sensitivity with respect to Z. While most engineering applications are interested in the convergence of moments of $\widetilde{U}(Z)$, i.e. $|E[\widetilde{U}(Z)^p] - E[U(Z)^p]| \to 0$, we focus our efforts on $L^p(\Omega)$ convergence which is stronger. Consequently, we propose refinement and stopping criterion for the adaptive method based on the L^p error and these criteria are quantitatively indicative of the performance and accuracy of $\widetilde{U}(Z)$ for a given computational budget. The surrogate model is constructed via various interpolants over Voronoi cells and we analyze the convergence properties of these interpolants.

In what follows, we first provide a review on collocation-based surrogates, namely stochastic collocation and the SROM-based surrogate. A comparison of the L^p error of these methods under varying stochastic dimension is presented via numerical examples. It is shown that stochastic collocation can underperform as it does not take into account the probability law of Z. We then elaborate on the proposed adaptive approach for SROM-based surrogates, address issues of implementation, and investigate its mathematical properties. The benefits of using the adaptive SROM-based surrogate, in which the probability law of Z is used to determine the location of the next SROM node, over the regular SROM-based surrogate is numerically demonstrated by comparing their L^p errors. Computationally, the two constructions differ in that the adaptive method has more evaluations of $U(Z)$, which is relatively inexpensive. The comparison was performed using examples with large stochastic dimension for test response functions and for a response derived from a stochastic PDE. Furthermore, we observe that the gap between the L^p errors of the adaptive and the regular SROM-based surrogate increases as p increases.

2 Collocation-based surrogates

We consider two types of collocation-based surrogates, namely stochastic collocation and the SROM-based surrogate. Both surrogates are constructed using a set of collocation nodes and interpolating functions. The convergence properties of these surrogates are examined and their construction and implementation are compared using numerical examples.

2.1 Stochastic collocation

Stochastic collocation method [9,10,19] constructs an interpolant $\widetilde{U}(Z)$ of $U(Z)$ over Γ. The process can be summarized in the following steps:

1. Select collocation points $\mathbf{z}_m \in \Gamma \subset \mathbb{R}^d$, $m = 1, \ldots, n$ and evaluate $U(\mathbf{z}_m)\, \forall m$.
2. Construct an interpolating function $\ell_m(Z)$ for each node $\mathbf{z}_m$ such that $\ell_m(\mathbf{z}_m) = 1$.
3. Formulate the surrogate as $\widetilde{U}_n(Z) = \sum_{m=1}^{n} U(\mathbf{z}_m)\ell_m(Z)$.

More specifically, suppose that $Z(\Omega) = \Gamma = [a_1, b_1] \times \cdots \times [a_d, b_d] \subset \mathbb{R}^d$. Denote by $Z^i = \{z_1^i, \ldots, z_{m_i}^i | z_k^i \in [a_i, b_i] \,\forall\, k\}$ the 1-dimensional set of collocation nodes in the i-th coordinate axis and let ℓ_k^i be the one-dimensional interpolating function corresponding to the node z_k^i, i.e. $\ell_k^i(z_k^i) = 1$. An interpolant for $U(Z)$ can be constructed as:

$$\widetilde{U}(Z) = \sum_{k_1=1}^{m_i} \cdots \sum_{k_d=1}^{m_d} \ell_{k_1}^1(z_{k_1}^1) \cdots \ell_{k_d}^d(z_{k_d}^d) \cdot U(z_{k_1}^1, \ldots, z_{k_d}^d) \tag{2.1}$$

which requires $m_1 \times \cdots \times m_d$ collocation nodes given by the Cartesian product $Z^1 \times \cdots \times Z^d$, i.e. the full grid.

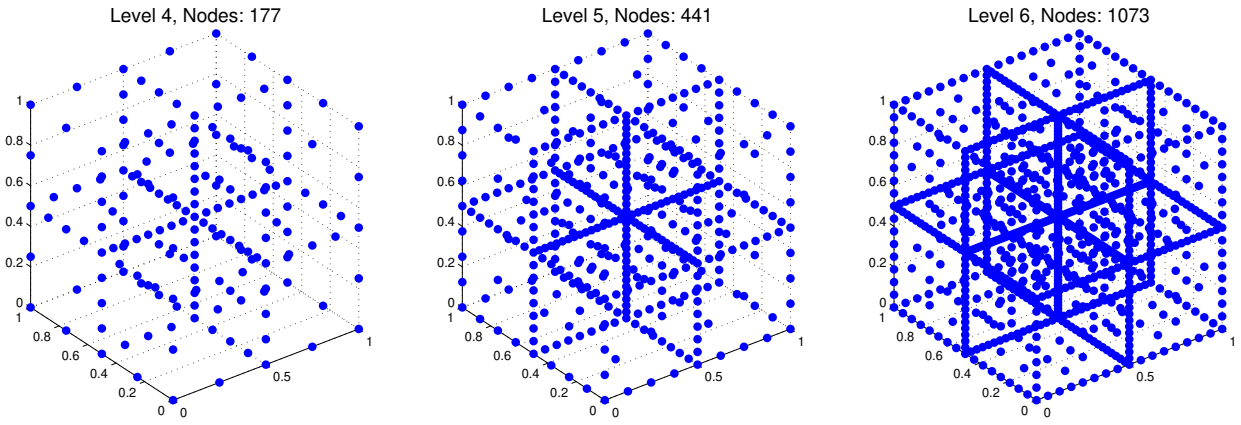

Figure 1: Illustration of sparse grid using Newton-Cotes equidistant nodes for $d = 3$ for different sparse grid levels. The nodes from a preceding level is a subset of the nodes in the succeeding level.

Because (2.1) suffers from the curse of dimensionality, sparse grids offer a more practical approach as the response is only evaluated at a subset of the full grid. An example of a 3-dimensional sparse grid using Newton-Cotes equidistant nodes in each coordinate axis is shown in Figure 1. The cardinality of this subset can be significantly lower than that of the full grid in high dimensions. The interpolant over this sparse grid is built through Smolyak's algorithm which involves taking the tensor product of one-dimensional interpolating functions in a special manner. Unlike (2.1), the sparse grid interpolant is constructed iteratively

from one sparse grid level to another but retains the same functional form. The differences are that the summation is now over the sparse grid nodes and that the weight of each basis function is the surplus between $U(Z)$ and the sparse grid interpolant in the preceding sparse grid level. An elaborate treatment of Smolyak's algorithm and sparse grid stochastic collocation is presented in [19] and the references therein.

If $\widetilde{U}$ is the sparse grid interpolant constructed using linear hat basis functions and nested Newton-Cotes equidistant nodes in each coordinate axis, the interpolation error is [19]:

$$\|U - \widetilde{U}\|_{L^\infty} = \mathcal{O}(N^{-2} \cdot |\log_2 N|^{3(d-1)})$$

where N is the number sparse grid nodes assuming that $U(z) \in C^2(\Gamma)$.

In 1-dimension, i.e. $d = 1$, the commonly used interpolating functions are the Lagrange basis functions and linear hat basis functions. For Lagrange interpolation, if $U \in C^n([a, b])$, $U^{(n+1)}(z)$ needs to be bounded on $[a, b]$ to guarantee uniform convergence. Because of this stringent regularity condition, the Lagrange interpolant can be prone to oscillatory behavior especially if the collocation points z_m are not dense at the boundary of Γ [16]. As such, collocation points such as Gauss-Legendre nodes or Clenshaw-Curtis nodes have to be employed. On the other hand, linear hat basis functions have local support and only require bounded second derivatives to ensure uniform convergence. Newton-Cotes equidistant nodes on $[a, b]$ (with endpoints included) is the most commonly used set of collocation points.

In summary, the key ingredients for stochastic collocation are the choice of interpolating functions and the choice of collocation nodes. The selection of collocation nodes is usually performed without regard for the probability law of Z. Recent works such as [7, 18] used Newton-Cotes equidistant nodes and utilized the probability density function of Z to determine where to add more collocation points in Γ. However, the resulting (sparse) grid is still structured and rectangular. To overcome these concerns, we introduce a different class of collocation-based interpolants that incorporate the probability law of Z into the construction of $\widetilde{U}(Z)$ and are well-suited for non-rectangular grids and for non-hypercube domains.

2.2 SROM-based surrogates

An SROM-based surrogate model $\widetilde{U}(Z)$ is a collocation-based surrogate for the stochastic response $U(Z)$. As with stochastic collocation, its construction entails the selection of collocation nodes and the selection of the interpolating function. We now elaborate how the collocation nodes can be chosen using the concept of an SROM.

Consider a random vector $Z \in \mathbb{R}^d$ defined on the probability space $(\Omega, \mathcal{F}, P)$. A Stochastic Reduced Order Model (SROM) of size m is a discrete random vector $\widetilde{Z}$ which takes on

values from the set $\{\widetilde{z}_k\}_{k=1}^m \subset \Gamma \subset \mathbb{R}^d$ with corresponding probabilities $\{p_k\}_{k=1}^m$. $\widetilde{Z}$ is defined on the same probability space as Z with $\{\widetilde{z}_k\}$ and $\{p_k\}$ chosen such that $\widetilde{Z}$ is statistically representative of Z. For instance, $\widetilde{Z}$ can be constructed such that the first two moments and the distribution of $\widetilde{Z}$ closely resemble that of Z. Different methods on the construction of SROMs are elaborated in [1, pp. 464-474] and the references therein. The non-uniqueness of the Stochastic Reduced Order Model is not an issue since we are only interested in obtaining samples which capture the statistics of the target random vector Z.

With the concept of an SROM for Z defined, an SROM-based surrogate model $\widetilde{U}(Z)$ can then be constructed using interpolants over Voronoi cells. In what follows, we consider Taylor-based interpolants and the Sibson's interpolant.

2.2.1 Taylor-based interpolants

The construction of an SROM-based surrogate model $\widetilde{U}(Z)$ using Taylor interpolants can be summarized in the following steps [2].

1. Generate a large number N of independent samples $\{z_k\}_{k=1}^N \subset \Gamma$ of Z.
2. For a specified m, $m \ll N$, construct an SROM $\widetilde{Z} = \{\widetilde{z}_k\}_{k=1}^m \subset \Gamma$ of size m for Z.
3. Partition the set $\{z_k\}_{k=1}^N$ into m subsets $\{\Gamma_j\}_{j=1}^m$ defined by $\Gamma_j = \{z \in \Gamma \ : \ \|z - \widetilde{z}_j\| < \|z - \widetilde{z}_l\| \text{ for } l = 1, \ldots, m,\ l \neq j\}$. If $\|z_k - \widetilde{z}_j\| = \|z_k - \widetilde{z}_l\|$, we can arbitrarily assign z_k to Γ_j or Γ_l. This step essentially performs a (discrete) Voronoi tessellation of Γ with the $\{\widetilde{z}_k\}_{k=1}^m$ as the centers. We therefore refer to Γ_j as Voronoi cells with $\cup_{j=1}^m \Gamma_j = \Gamma$.
4. On each Voronoi cell Γ_j, perform a zeroth-order or first-order Taylor expansion of Z at $\widetilde{z}_j$ which gives rise to

$$\widetilde{U}_m(Z) = \sum_{k=1}^m \mathbb{1}_{(Z \in \Gamma_k)} U(\widetilde{z}_k) \tag{2.2}$$

or

$$\widetilde{U}_m(Z) = \sum_{k=1}^m \mathbb{1}_{(Z \in \Gamma_k)} [U(\widetilde{z}_k) + \nabla U(\widetilde{z}_k) \cdot (Z - \widetilde{z}_k)], \tag{2.3}$$

respectively, where $\nabla U(\widetilde{z}_k)$ denotes the gradient of $U(z)$ at $z = \widetilde{z}_k$.

An illustration of the resulting surrogate model is presented in Example 1 and Figure 2.

Example 1. Let $Z = (Z_1, Z_2)$ where $Z_1 \sim F_1^{-1}(\Phi(Y_1)), Z_2 \sim F_2^{-1}(\Phi(Y_2))$ with Φ being the standard normal CDF, $Y_1, Y_2 \sim N(0,1), \text{cov}(Y_1, Y_2) = 0.3$ and F_1, F_2 are the CDFs of beta$(2,6)$ and beta$(6,6)$, respectively. Suppose that we want to construct a surrogate for $U(Z) = \sin(Z_1 + 3(Z_2 - 0.2))$. The left plot of Figure 2 shows an SROM $\widetilde{Z}$ of size $m = 15$

and the right plot exhibits the response function $U(Z)$ and the surrogate model. We remark that the nodes of $\widetilde{Z}$ are concentrated in regions of Z with high probability.

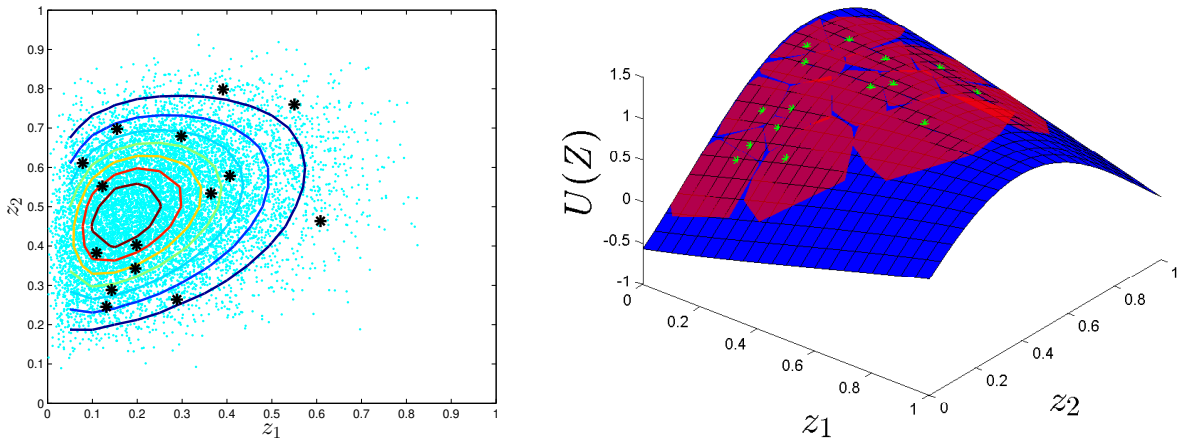

Figure 2: Left – SROM $\widetilde{Z}$ (black asterisks) for Z in Example 1 with samples of Z (blue dots) and contour of the joint pdf of Z. Right – Illustration of the surrogate $\widetilde{U}_m(Z)$ (red surface) for $U(Z)$ (blue surface) using first-order Taylor interpolants for $m = 15$ where Z is defined as in Example 1. The green nodes are the SROM nodes.

In sum, an SROM-based surrogate as developed in [2] uses an SROM of size m for Z as its collocation nodes while the interpolating functions are piecewise constant functions in each Voronoi cell as in (2.2) or, as shown in Figure 2 (right), piecewise tangent hyperplanes in each Voronoi cell as in (2.3). The interpolating functions are chosen as such in order to accommodate the fact that the collocation points no longer lie on a rectangular grid. We note that higher-order Taylor expansions can be used as interpolants over each Voronoi cell. However, this would require stringent regularity conditions of $U(Z)$ to guarantee convergence of $\widetilde{U}_m(Z)$ as we shall see below. The computational cost would also scale non-linearly to compute the higher-ordered partial derivatives, making it impractical in high dimensions.

We now briefly remark on some computational aspects involving the SROM-based surrogate and subsequently review its important properties as presented in [2]. We focus on the SROM-based surrogate using 1st-order Taylor interpolants because the surrogate in (2.2) is a coarse approximation as it does not adequately capture the variation of $U(Z)$. Hence, $\widetilde{U}_m(Z)$ in what follows refers to (2.3). Using $\widetilde{U}_m(Z)$, statistical information of $U(Z)$ can now be approximated at a cheaper computational cost. For example, the p-th order moments of $U(Z)$ can be estimated as follows [1, pp. 470]:

$$E[U(Z)^p] \simeq E[\widetilde{U}_m(Z)^p] \simeq \sum_{k=1}^{m} \frac{n_k}{N} \sum_{z_j \in \Gamma_k} \frac{1}{n_k} (U(\widetilde{z}_k) + \nabla U(\widetilde{z}_k) \cdot (z_j - \widetilde{z}_k))^p$$

where n_k is the cardinality of the set $\{j | z_j \in \Gamma_k\}$. Evidently, it is substantially cheaper to evaluate $U(\widetilde{z}_k) + \nabla U(\widetilde{z}_k) \cdot (z_j - \widetilde{z}_k)$ compared to the true response. The fact that $E[U(Z)^p]$ can be estimated by $E[\widetilde{U}_m(Z)^p]$ can be justified through the properties below.

Theorem 1. *Let $\widetilde{U}_m(Z)$ be as in (2.3). Suppose that $Z(\Omega) = \Gamma \subset \mathbb{R}^d$ is an open convex domain such that $U(Z) \in C^2(\Gamma)$. Furthermore, assume that the second-order partial derivatives of $U(Z)$ are bounded, i.e. $\left| \frac{\partial^2 U(z)}{\partial z^{(r)} \partial z^{(s)}} \right| \leq M$ for all $r, s = 1, \ldots, d, z \in \Gamma$. For $z \in \Gamma_k$, where Γ_k is a Voronoi cell of Γ with center $\widetilde{z}_k$, we then have that for some constant C_k*

$$|U(z) - \widetilde{U}_m(z)| \leq \frac{C_k}{2} \|z - \widetilde{z}_k\|^2. \tag{2.4}$$

Proof. Let $z \in \Gamma_k$. By Taylor's remainder theorem, cf. [28],

$$U(z) - \widetilde{U}_m(z) = \frac{1}{2} (z - \widetilde{z}_k)^T \; H(U(\widetilde{\xi}_k)) \; (z - \widetilde{z}_k)$$

where $\widetilde{\xi}_k$ is a point on the line segment joining z and $\widetilde{z}_k$ and $H(U(\widetilde{\xi}_k))$ is the Hessian matrix of $U(z)$ evaluated at $z = \widetilde{\xi}_k$. The symmetry of the Hessian matrix implies that $H(U(\widetilde{\xi}_k)) = Q^T \Lambda Q$ where Q is an orthogonal matrix, Λ is a diagonal matrix with entries $\lambda_i(\widetilde{\xi}_k)$, and that Q, Λ depend on $\widetilde{\xi}_k$. Hence, if $y := Q \cdot (z - \widetilde{z}_k)$,

$$\begin{aligned} |U(z) - \widetilde{U}_m(z)| &= \frac{1}{2} |y^T \Lambda y| = \frac{1}{2} \left| \sum_{i=1}^{d} \lambda_i(\widetilde{\xi}_k) \, y_i^2 \right| \\ &\leq \frac{1}{2} \left(\max_i |\lambda_i(\widetilde{\xi}_k)| \right) \|y\|^2 \leq \frac{C_k}{2} \|z - \widetilde{z}_k\|^2 \end{aligned} \tag{2.5}$$

where C_k is a bound on $\max_i |\lambda_i(\widetilde{\xi}_k)|$ owing to the fact that the entries of the Hessian matrix are bounded.

□

Using Theorem 1, we can easily prove convergence properties of the SROM-based surrogate.

Corollary 2. *Suppose that the second-order partial derivatives of $U(Z)$ are bounded. Consider a refining sequence of SROMs for Z, that is, if $\widetilde{Z}_m = \{\widetilde{z}_k\}_{k=1}^m$ and $\widetilde{Z}_{m+1} = \{\widetilde{z}_k\}_{k=1}^{m+1}$ are SROMS for Z, $\widetilde{Z}_m \subset \widetilde{Z}_{m+1}$. We then have that $\widetilde{U}_m(Z) \to U(Z)$ almost surely as $m \to \infty$. In addition, if $\exists\, p \geq 1$ such that $U(Z) \in L^p(\Omega)$ then $\widetilde{U}_m(Z)$ also converges to $U(Z)$ in $L^p(\Omega)$ [2].*

Proof. See [2], p. 273.

□

Corollary 2 illustrates how the SROM-based surrogate (2.3) possesses strong convergence properties under very mild conditions. Because a.s. and L^p convergence each imply convergence in distribution, we have that for every $u \in \mathbb{R}$, $P(\widetilde{U}_m(Z) \leq u) \to P(U(Z) \leq u)$ as $m \to \infty$. In addition, convergence in L^p implies convergence in p^{th} order moments, i.e. $E[\widetilde{U}_m(Z)^p] \to E[U(Z)^p]$, which is often of interest in engineering applications. Hence, the convergence properties of (2.3) are comparable to that of the sparse grid stochastic collocation using linear hat basis and Newton-Cotes equidistant nodes under identical regularity conditions on $U(Z)$.

2.2.2 Sibson-based interpolant

We extend the class of interpolants introduced in [2] for the SROM-based surrogate by considering an interpolant from geoscience. Subsequently, we will state and prove properties of this interpolant and comment on its implementation.

Suppose that we have at our disposal $\{(\widetilde{z}_k, U(\widetilde{z}_k))\}_{k=1}^m$ where $\widetilde{z}_k \in \mathbb{R}^d$ and that we are concerned with approximating the mapping $Z \mapsto U(Z)$ on $\Gamma \subset \mathbb{R}^d$ where Γ is the interior of the convex hull of the nodes $\{\widetilde{z}_k\}_{k=1}^m$. This can be accomplished using Sibson's interpolation [22] in which the value of the interpolant at a new point $z^* \in \Gamma$ is obtained as a weighted sum of $U(\widetilde{z}_k)$ for $\widetilde{z}_k$ close to z^*. The weights are computed by determining how much of the volume in each Voronoi cell of the original tessellation generated using $\{\widetilde{z}_k\}_{k=1}^m$ as the centers is lost upon retessellation with $z^* \cup \{\widetilde{z}_k\}_{k=1}^m$ as the new centers.

More formally, assume that we have a Voronoi tessellation of $\mathbb{R}^d$, not Γ, denoted by $\{\Gamma_k\}_{k=1}^m$ with $\{\widetilde{z}_k\}_{k=1}^m$ as the centers. Two Voronoi cells Γ_i and Γ_j are considered neighbors if Γ_i and Γ_j have a common boundary. For a given $z^* \in \Gamma$ the value of $U(z^*)$ can be approximated by a surrogate $\widetilde{U}_m(z^*)$ as follows:

1. If $z^* \in \{\widetilde{z}_k\}_{k=1}^m$ then $\widetilde{U}_m(z^*) = U(\widetilde{z}_k)$ for some $k \leq m$.
2. Otherwise, construct a Voronoi tessellation $\{\widetilde{\Gamma}_k\}_{k=1}^m \cup \{\widetilde{\Gamma}^*\}$ of $\mathbb{R}^d$ with $\{\widetilde{z}_k\}_{k=1}^m \cup \{z^*\}$ as the corresponding centers.

3. For each $k = 1, \ldots, m$, compute

$$w_k(z^*) := \frac{\lambda(\Gamma_k \cap \widetilde{\Gamma}^*)}{\lambda(\widetilde{\Gamma}^*)} \tag{2.6}$$

where $\lambda(\cdot)$ is the Lebesgue measure. The numerator of (2.6) represents how much volume from Γ_k is lost and absorbed into $\widetilde{\Gamma}^*$ upon the inclusion of z^* as a Voronoi center. We remark that $\lambda(\Gamma_k \cap \widetilde{\Gamma}^*)$ and $\lambda(\widetilde{\Gamma}^*)$ are well-defined because $\widetilde{\Gamma}^*$ is a bounded region owing to the fact that z^* is in the interior of the convex hull.

4. Set $\widetilde{U}_m(z^*) = \sum_{k=1}^{m} w_k(z^*) U(\widetilde{z}_k)$ to be the value of the surrogate at z^*.

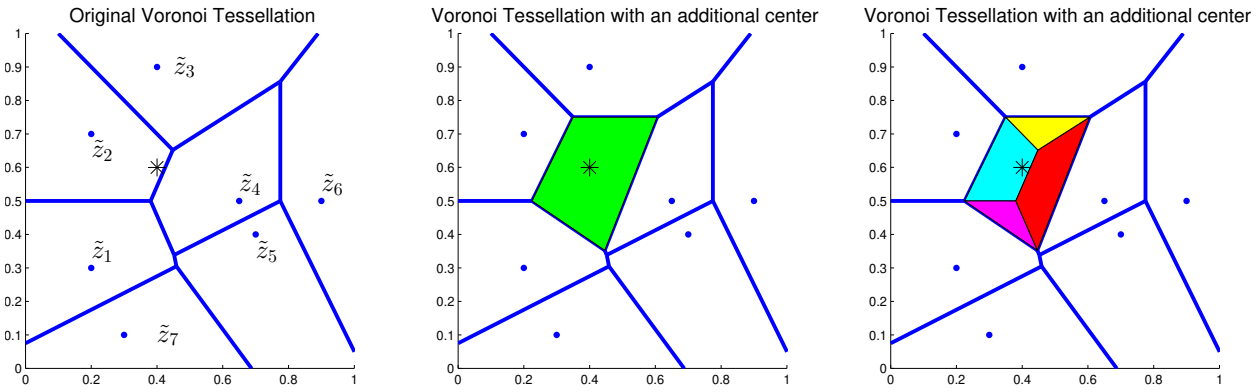

Figure 3: Illustration of the computation of the weights in Sibson's interpolant.

An example of how the weights in Sibson's interpolant are computed is illustrated in Figure 3. The nodes $\{\widetilde{z}_k\}$ are denoted by the blue dots with their corresponding Voronoi cells in the leftmost plot while the asterisk represents z^*. The updated Voronoi tessellation after inserting z^* as a Voronoi center is shown in the middle plot which corresponds to step 2 in the above procedure. The weights $w_k(z^*)$ for $\widetilde{z}_k$, $k = 1, \ldots, 4$, can be computed as the ratio of the area of the respective colored regions $(\Gamma_k \cap \widetilde{\Gamma}^*)$ in the rightmost plot with the area of the green colored region $(\widetilde{\Gamma}^*)$ in the middle plot. Furthermore, $w_k(z^*) = 0$ for $k \geq 5$ which demonstrates that Sibson's interpolant only uses the values at the neighboring nodes.

Sibson's interpolant possesses a number of properties that have been stated and proven in [23–25]. We review some of them here.

Property 3 (Sibson 1980)**.** *The functions $w_k(z^*)$ satisfy the local coordinates property in that*

$$z^* = \sum_{k=1}^{m} w_k(z^*) \widetilde{z}_k$$

for all $z^* \in \Gamma \subset \mathbb{R}^d$.

Property 4. *The interpolant $\widetilde{U}_m(z^*)$ has linear precision. In other words, if the true response U is linear, $U(z^*) = \widetilde{U}_m(z^*)\ \forall z^* \in \Gamma$.*

The second property easily follows from the local coordinates property. Indeed, if U is linear,

$$U(z^*) = U\left(\sum_{k=1}^{m} w_k(z^*)\widetilde{z}_k\right) = \sum_{k=1}^{m} w_k(z^*)U(\widetilde{z}_k) = \widetilde{U}_m(z^*).$$

Property 5. *$\widetilde{U}_m(z^*)$ is continuous for all $z^* \in \Gamma$ and is continuously differentiable for all $z^* \in \Gamma$, $z^* \notin \{\widetilde{z}_k\}_{k=1}^m$.*

Property 6. *For $d = 1$, i.e. in the 1-dimensional case, Sibson's interpolant and collocation with linear hat basis functions coincide. For $d = 2$, it has been shown in [24] that Sibson's interpolant has gridded bilinear precision. In other words, suppose that $\{\widetilde{z}_k\}_{k=1}^m$ lie on a rectangular lattice, i.e. $\{\widetilde{z}_k\}_{k=1}^m$ are the Newton-Cotes equidistant nodes in $\Gamma \subset \mathbb{R}^2$. Sibson's interpolant then coincides with collocation using $\{\widetilde{z}_k\}_{k=1}^m$ and linear hat basis functions.*

Based on the above property, Sibson's interpolant is a generalization of collocation using linear hat basis functions for arbitrarily located collocation nodes, for any dimension. Hence, using this interpolant, the nodes can be chosen in accordance with the probability law of Z. However, the properties above are only valid in the interior of the convex hull of $\{\widetilde{z}_k\}_{k=1}^m$. Efforts to extend the validity of the interpolant beyond the interior have been undertaken as in [26, 27]. These approaches introduce a different or additional geometric construction and we do not pursue them here for simplicity.

It is now straightforward to incorporate Sibson's interpolant into the SROM-based surrogate model – we simply let the $\{\widetilde{z}_k\}_{k=1}^m$ above to be an SROM for the random vector Z. Subsequently, we analyze the convergence properties of Sibson's interpolant when employed to construct surrogate models for a stochastic response $U(Z)$. In the following, we will assume that Γ is the interior of the convex hull of $\{\widetilde{z}_k\}_{k=1}^m$.

Proposition 7. *Suppose that $U(Z)$ is continuous on Γ. Consider a refining sequence of SROMs for Z. We then have $\widetilde{U}_m(Z) \to U(Z)$ as $m \to \infty$.*

Proof. Let $\epsilon > 0$ be given and let $\{\widetilde{z}_k\}_{k=1}^m \subset \Gamma$ be the SROM nodes which partition Γ into the corresponding Voronoi cells $\{\Gamma_k\}_{k=1}^m$. For $z^* \in \Gamma$, we can always choose m to be large enough so that $|U(\widetilde{z}_k) - U(z^*)| < \epsilon,\ \forall\, k$ such that $w_k(z^*) \neq 0$ due to the continuity of $U(Z)$ at $Z = z^*$. By noting that $\sum_{k=1}^{m} w_k(z^*) = 1$, the discrepancy can be bounded by

$$\left|\widetilde{U}_m(z^*) - U(z^*)\right| = \left|\sum_{k=1}^{m} w_k(z^*)U(\widetilde{z}_k) - U(z^*)\right| = \left|\sum_{k=1}^{m} w_k(z^*)(U(\widetilde{z}_k) - U(z^*))\right|$$

$$\leq \sum_{k=1}^{m} w_k(z^*)|U(\widetilde{z}_k) - U(z^*)| < \epsilon.$$

□

Proposition 8. *Suppose that $U(Z)$ is differentiable on Γ and that $\left|\frac{\partial U(z)}{\partial z^{(i)}}\right| \leq M$ for $i = 1, \ldots, d$ and $z \in \Gamma$. Consider a refining sequence of SROMs for Z. We then have that for $p \geq 1$, if $U(Z) \in L^p(\Omega)$, $\widetilde{U}_m(Z) \to U(Z)$ in $L^p(\Omega)$ as $m \to \infty$.*

Proof. By the Mean Value Theorem,

$$|U(\widetilde{z}_k) - U(z)| \leq \|\nabla U(\xi)\| \cdot \|\widetilde{z}_k - z\| \leq \sqrt{d}M\|\widetilde{z}_k - z\|$$

where ξ lies on the line segment joining z and $\widetilde{z}_k$. As a consequence,

$$\begin{aligned}\left|\widetilde{U}_m(z) - U(z)\right| &\leq \sum_{k=1}^{m} w_k(z)|U(\widetilde{z}_k) - U(z)| \leq \sqrt{d}M \max_{k,\, w_k \neq 0} \|\widetilde{z}_k - z\| \\ &\leq \sqrt{d}M \max_{z\in\Gamma} \max_{k,\, w_k \neq 0} \|\widetilde{z}_k - z\|\end{aligned}$$

where $\max\limits_{z\in\Gamma} \max\limits_{k,\, w_k\neq 0} \|\widetilde{z}_k - z\| \to 0$ as $m \to \infty$, i.e. for $\epsilon > 0$, $\exists N$ such that $\max\limits_{z\in\Gamma} \max\limits_{k,\, w_k\neq 0} \|\widetilde{z}_k - z\| < \epsilon$ for $m \geq N$. The claim now holds true if $p = \infty$. On the other hand,

$$\|\widetilde{U}_m(Z) - U(Z)\|^p_{L^p(\Omega)} = \int_\Gamma |\widetilde{U}_m(z) - U(z)|^p \, dF(z) < (\sqrt{d}M\epsilon)^p.$$

□

Remark. The proof above demonstrates that the L^p error of the SROM-based Sibson's interpolant is bounded by the size of the Voronoi cell as in the case of the SROM-based surrogate models introduced earlier. It is also straightforward to show that the conditions on differentiability and bounded derivatives of $U(Z)$ in Proposition 8 can be relaxed by assuming that $U(Z)$ is uniformly continuous on Γ.

In addition to its convergence properties, the SROM-based Sibson's interpolant guarantees that the surrogate does not attain unrealistic values:

Property 9. *If $U_{min} = \min\limits_{z\in\Gamma} U(z)$ and $U_{max} = \max\limits_{z\in\Gamma} U(z)$, then $P(\widetilde{U}_m(Z) \in [U_{min}, U_{max}]) = 1$.*

This property easily follows from the fact that $\sum_{k=1}^{m} w_k(z) = 1$.

This is important in applications in which it is preferable that the surrogate respects the physics of the model aside from guaranteeing convergence.

Despite the flexibility offered by Sibson's interpolant, its limitations are clear: the computation of the weights w_k poses difficulties in high stochastic dimension, a challenge that is not encountered in the 2-d or 3-d applications in computational geometry, geoscience, etc. For one, volumes of Voronoi cells need to be approximated. Since determining the exact boundaries of Voronoi cells might be impractical in high dimensions, a counting process can instead be utilized to estimate volumes through Monte Carlo integration. Secondly, for z^* close to the boundary of the convex hull of $\{\widetilde{z}_k\}_{k=1}^m$, the boundaries of the resulting Voronoi cell $\widetilde{\Gamma}^*$ can extend well beyond the interior of the convex hull Γ even though its Lebesgue measure is guaranteed to be finite. Consequently, a "bounding box" encompassing Γ is used in applications to impose bounds on the extent of $\widetilde{\Gamma}^*$, making it more feasible to estimate its volume. The accuracy of this interpolant therefore depends on the size of the bounding box as well as the number of samples for the counting process.

2.3 Comparison of collocation-based surrogates

After providing a survey of collocation-based surrogates, a comparison of their key features is summarized in Table 2.1.

	SROM-based			Collocation (linear hat)	
	0th-order Taylor	1st-order Taylor	Sibson	Full grid	Sparse grid
Gradient needed?	✗	✓	✗	✗	✗
Exact for linear response?	✗	✓	✓	✓	✓
Is $\widetilde{U}_m(z)$ within bounds?	✓	✗	✓	✓	✗

Table 2.1: Comparison of collocation-based surrogates.

In this section, we are primarily interested in investigating the effect of the choice of collocation nodes through a comparison of the SROM-based surrogate with sparse grid stochastic collocation for varying stochastic dimension and for the same computational budget. The

1st-order Taylor interpolant will be used for the SROM-based surrogate while Newton-Cotes equidistant nodes and linear hat basis function will be used for sparse grid stochastic collocation. We have chosen to focus on the 1st-order Taylor interpolant because it provides a balance between accuracy and ease of implementation. The specification of the response is outlined in the following example.

Example 2. Consider the stochastic response $U(Z) = \arctan(15 \cdot \|Z - 0.7\|^2) - \arctan(15 \cdot \|Z - 0.3\|^2)$, $Z \in \mathbb{R}^d$, where $Z_i \sim F^{-1}(\Phi(Y_i))$ with Φ being the standard normal CDF, $Y_i \sim N(0,1), \mathrm{cov}(Y_i, Y_j) = 0.85$ for $i \neq j$ and F_i is the CDF of beta$(3,3)$. An illustration of the probability law of Z and a plot of $U(z)$ for $z \in [0,1] \times [0,1]$ is shown in Figure 4 for $d = 2$.

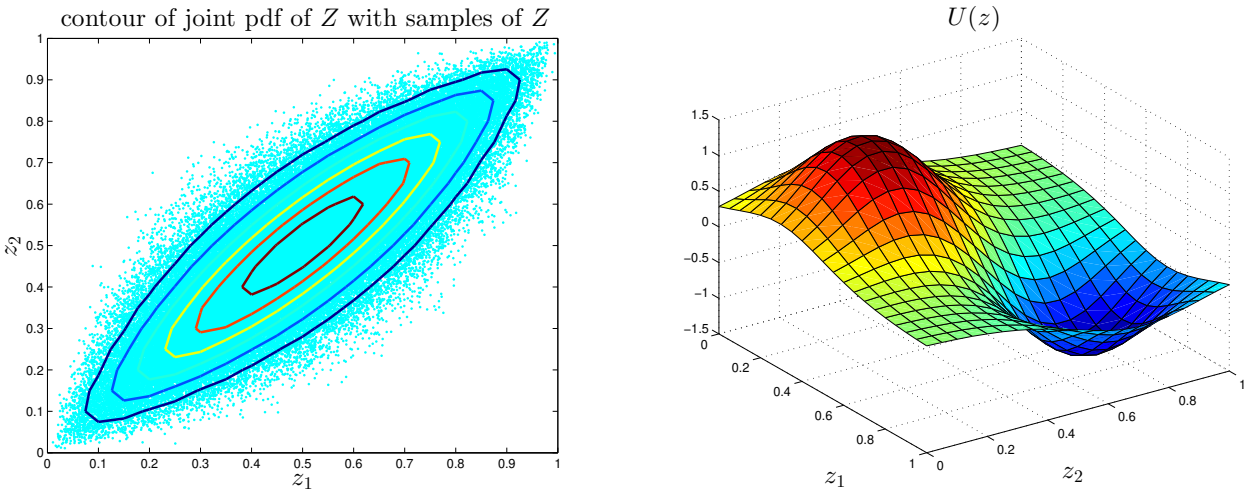

Figure 4: Joint pdf of Z (left) and $U(z)$ (right).

Figures 5, 6, 7, exhibit the L^p errors $\|U(Z) - \widetilde{U}(Z)\|_{L^p(\Omega)}$ for $d = 2, 4, 6$, respectively, obtained via Monte Carlo approximations. For each figure, the blue solid curve (denoted by SpGrid) and the red dashed curve (denoted by SROM) are the L^p errors for the sparse grid surrogate and the SROM-based surrogate, respectively. Each figure contains 4 subplots which correspond to $p = 1, 2, 3, \infty$. The x-axis of each subplot refers to the sparse grid level which is characterized by the number of evaluations of $U(z)$ required to form the interpolant. Hence, for each sparse grid level with n_l nodes, $\lfloor n_l/(d+1) \rfloor$ SROM nodes are used to ensure that the number of computational units between both surrogates is similar.

As the L^p error comparison plots demonstrate for this example, sparse grid stochastic collocation outperforms SROM-based surrogate for $d = 2$ even though the sparse grid nodes do not align with the high probability regions of Z. In Figure 8, we show the location of the sparse grid nodes for the 8^{th} sparse grid level (characterized by the location of the sparse grid nodes as in Figure 1) and the corresponding SROM nodes for the same computational budget, together with the contour of the response. It is therefore not surprising that the

sparse grid collocation perform well because the nodes are adequately distributed in the domain. However, this trend is reversed as the stochastic dimension is increased, as can be seen in Figures 6 and 7. For skewed distributions for Z in high dimensions, a large sparse grid level, which implies a large computational expense, is necessary for the sparse grid nodes to capture the most probable regions of Z. This numerical example thus underscores the importance of incorporating the probability law of Z into the location of the collocation nodes.

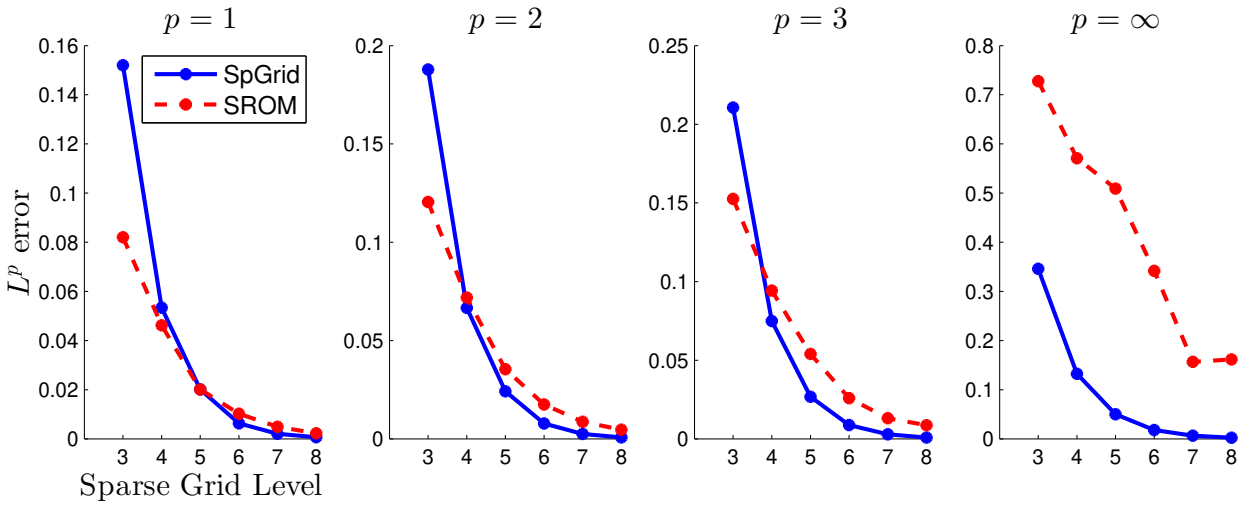

Figure 5: L^p error comparison of SROM-based surrogate and sparse grid stochastic collocation for $d = 2$. Sparse grid level/Number of Nodes: 4/65, 6/321, 8/1537.

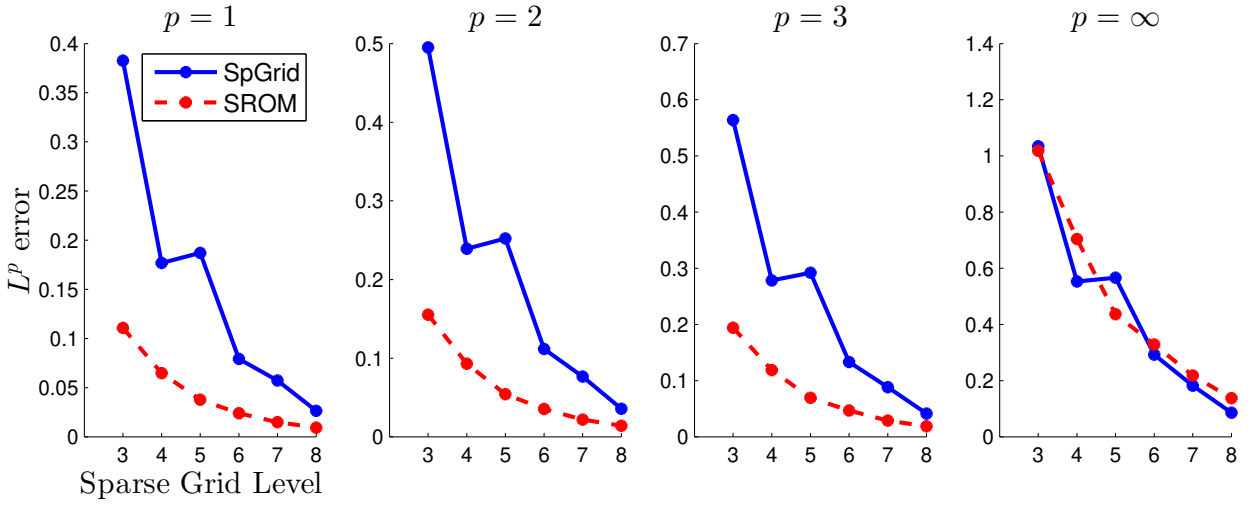

Figure 6: L^p error comparison of SROM-based surrogate and sparse grid stochastic collocation for $d = 4$. Sparse grid level/Number of Nodes: 4/401, 6/2929, 8/18945.

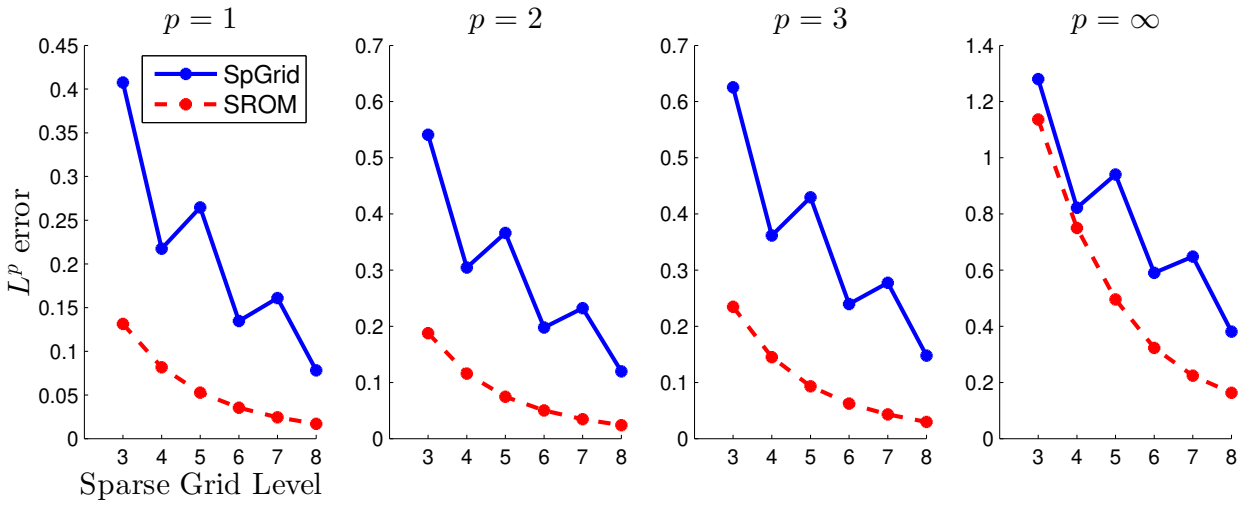

Figure 7: L^p error comparison of SROM-based surrogate and sparse grid stochastic collocation for $d = 6$. Sparse grid level/Number of Nodes: 4/1457, 6/15121, 8/127105.

3 Adaptive SROM-based surrogate

We now focus on the main contribution of this work. While the SROM-based surrogate is capable of capturing the probability law of Z, we can further extend this surrogate in order to also capture the regions of Z for which $U(Z)$ manifests substantial variation. This can be achieved by constructing the SROM-based surrogate in a sequential manner. We offer two approaches for the adaptive construction wherein one employs a global sampling strategy while the other uses a local sampling strategy. The algorithm presented below which incorporates global sampling has some similarities with an algorithm coined for the deterministic case in a terrain modeling application [29], although both algorithms have been developed independently.

3.1 Description and illustration of the algorithm

As before, we will use the SROM-based surrogate with 1st-order Taylor interpolant due to reasons specified above. The main idea behind the adaptive algorithm is that instead of simultaneously computing the gradients at all Voronoi centers, the centers at which the gradients will be evaluated will be chosen sequentially. Therefore, we will assume that the following parameters are specified:

- p – the order of the L^p error ($p \geq 1$).
- ϵ – tolerance level for the L^p error.

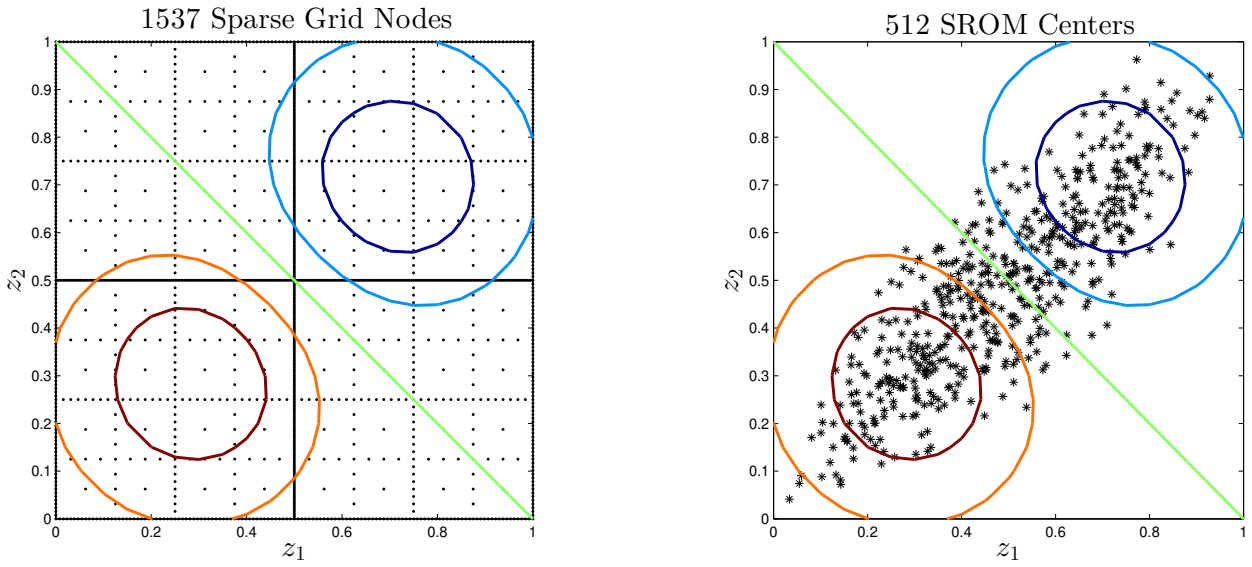

Figure 8: Location of sparse grid nodes for the 8th level (left) and SROM nodes (right) with contours of $U(z)$.

- τ – the available computational units/budget.

3.1.1 Adaptive method with global sampling

The global sampling approach proceeds with the initialization step in which we generate an SROM $\widetilde{Z} = \{\widetilde{z}_k\}_{k=1}^M$ of Z and evaluate $U(\widetilde{z}_k)\,\forall k$, but not $\nabla U(\widetilde{z}_k)$, within the available computational budget τ. In order to obtain an initial surrogate model for $U(z)$, we then construct an SROM $\widetilde{Z}^*$ of $\widetilde{Z}$ of size $m \ll M$ (i.e. $m = 1$), where the SROM nodes of $\widetilde{Z}^*$ is a subset of $\{\widetilde{z}_k\}_{k=1}^M$. The initial surrogate model (2.3) can then be constructed using the SROM nodes $\widetilde{Z}^*$ which induces an initial partitioning of Γ through $\{\Gamma_i\}_{i=1}^m$. At every iteration, we select a partition where the error of the surrogate is largest and select an SROM node $\widetilde{z}_k$ from this partition under some criteria. The gradient at this SROM node is then computed, a new partition corresponding to this SROM node is then constructed, and the current surrogate is updated owing to the new partition. This procedure is repeated until the computational budget τ is exhausted or if the error is smaller than the specified tolerance ϵ.

The construction of the adaptive SROM-based surrogate with global sampling can thus be summarized as follows:

Algorithm 1 Adaptive method for SROM-based surrogate with global sampling

1: Generate an SROM $\widetilde{Z}$ of Z given by $\mathcal{A} = \{\widetilde{z}_k\}_{k=1}^M$ and evaluate $U(\widetilde{z}_k)$.
2: Generate an SROM $\widetilde{Z}^*$ of $\widetilde{Z}$ given by $\{\widetilde{z}_k^*\}_{k=1}^m \subset \mathcal{A}$, $m \ll M$, and evaluate $\nabla U(\widetilde{z}_k^*)$.
3: Construct a Voronoi tessellation $\{\Gamma_i\}_{i=1}^m$ of Γ using $\{\widetilde{z}_k^*\}_{k=1}^m$ as the centers.
4: Compute the surrogate $\widetilde{U}_m(z)$ and update $\mathcal{A} \leftarrow \mathcal{A} \setminus \{\widetilde{z}_k^*\}_{k=1}^m$.
5: **while** budget τ is not exhausted **and** approximate $\|\widetilde{U}_m(z) - U(z)\|_{L^p(\Gamma)}^p > \epsilon$ **do**
6: Select *partition* $\Gamma_k \in \{\Gamma_i\}_{i=1}^m$ with the largest approximate $L^p(\Gamma_k)$ error.
7: *Select node* $\widetilde{z}_j \in \Gamma_k \cap \mathcal{A}$ and compute $\nabla U(\widetilde{z}_j)$.
8: *Refine the partition* Γ_k.
9: Update the surrogate $\widetilde{U}_{m+1}(z) \leftarrow \widetilde{U}_m(z)$, the set $\mathcal{A} \leftarrow \mathcal{A} \setminus \{\widetilde{z}_j\}$, and $m \leftarrow m + 1$.
10: **end while**

We now address the technicalities encountered in the algorithm above. At each iteration, the L^p error of the surrogate is used as the refinement measure to drive adaptivity. Mathematically, the L^p error of $\widetilde{U}_m(z)$ on the *partition* Γ_k is equivalent to

$$\begin{aligned}\|U(z) - \widetilde{U}_m(z)\|_{L^p(\Gamma_k)}^p &= \int_{\Gamma_k} |U(z) - \widetilde{U}_m(z)|^p \, dF(z) \\ &= E[|U(z) - \widetilde{U}_m(z)|^p \mid Z \in \Gamma_k] \cdot P(\Gamma_k),\end{aligned}$$

i.e., the refinement measure of Γ_k is equal to the average p-th order surplus of the surrogate weighted by the probability of that *partition*. As the refinement measure cannot be computed exactly in practice, we approximate $P(\Gamma_k)$ by determining the proportion of samples of Z such that $z \in \Gamma_k$ whereas the average p-th order surplus conditional on Γ_k can be approximated using the response values $U(\widetilde{z}_j)$ for $\widetilde{z}_j \in \Gamma_k \cap \mathcal{A}$ where $\{\widetilde{z}_j\}$ is obtained from the initialization step above. More explicitly, if we denote by $n_k \coloneqq \#\{\widetilde{z}_j | \widetilde{z}_j \in \Gamma_k \cap \mathcal{A}\}$,

$$E[|U(z) - \widetilde{U}_m(z)|^p \mid Z \in \Gamma_k] \approx \frac{1}{n_k} \sum_{i=1}^{n_k} |U(\widetilde{z}_i) - \widetilde{U}_m(\widetilde{z}_i)|^p.$$

As the iteration of the adaptive algorithm with global sampling progresses, it is possible that $\{\widetilde{z}_j | \widetilde{z}_j \in \Gamma_j \cap \mathcal{A}\}$ may eventually be empty for some partition Γ_j and that the L^p error may not be approximated. To preempt such a scenario, we will assume that when the user specifies τ computational units in the initialization stage, this already includes n_c units of contingency response evaluations of $U(z)$, i.e. $\tau = M + g + n_c$ where g is the maximum number of gradient calculations. If this situation occurs, we construct an SROM of size $\lceil P(\Gamma_j) n_c \rceil$ for Z conditioned on Γ_j and evaluate $U(z)$ on these newly obtained SROM nodes to approximate the L^p error on this *partition*. Subsequently, we add these SROM nodes to $\mathcal{A}$ and we update n_c as $n_c \leftarrow n_c - \lceil P(\Gamma_j) n_c \rceil$.

We can also obtain an approximate confidence interval to quantify the error in our estimation of the refinement measure of Γ_k. Let $Y_i = |U(\widetilde{z}_i) - \widetilde{U}_m(\widetilde{z}_i)|^p$, $i = 1, \ldots, n_k$. By the Central Limit Theorem, $S_{n_k} = \dfrac{1}{n_k}\displaystyle\sum_{i=1}^{n_k} Y_i$ is asymptotically $\mathcal{N}\left(\mu, \dfrac{\sigma^2}{n_k}\right)$ where $\mu = E[|U(z) - \widetilde{U}_m(z)|^p \,|\, Z \in \Gamma_k]$ while $\sigma^2 \approx \dfrac{1}{n_k - 1}\displaystyle\sum_{i=1}^{n_k}(Y_i - S_{n_k})^2$. Confidence bands for μ are now immediate from this asymptotic distribution.

The choice of refinement measure we have pursued implies that the L^p error over Γ will be used as a stopping criterion for the adaptive process. The same techniques described above are applicable in approximating the L^p error on the whole probability space. Thus, the adaptive process terminates either when the approximate L^p norm is less than the specified tolerance or when the budget τ is exhausted. The latter occurs precisely when the maximum number of gradient evaluations has been reached or when there are no more contingency response evaluations available.

3.1.2 Adaptive method with local sampling

An alternative to the adaptive algorithm with global sampling as presented above is possible through a local sampling procedure. Essentially, instead of selecting all of the samples $\widetilde{z}_k$ at which to evaluate the response at the beginning of the algorithm, some of these samples are selected as the adaptive construction progresses. The initialization step is still the same, however, rather than generating an SROM $\widetilde{Z}$ of size M as in before, we only generate an SROM $\widetilde{Z} = \{\widetilde{z}_k\}_{k=1}^{M'}$ where $M' \ll M$ and evaluate $U(\widetilde{z})\,\forall k$ but not $\nabla U(\widetilde{z}_k)$. Subsequently, at each iteration, for each partition Γ_i, we construct an SROM of size m_i for Z conditioned on Γ_i and evaluate $U(z)$ on these SROM nodes. This local sampling scheme removes the need to determine the location of all nodes $\widetilde{z}_k$ at the initialization step.

The construction of the adaptive SROM-based surrogate with local sampling is summarized below:

Algorithm 2 Adaptive method for SROM-based surrogate with local sampling

1: Perform Steps 1-4 as in Algorithm 1.
2: **while** budget τ is not exhausted **and** approximate $\|\widetilde{U}_m(z) - U(z)\|^p_{L^p(\Gamma)} > \epsilon$ **do**
3: Perform Lines 6-9 as in Algorithm 1.
4: **for** each Γ_i **do**
5: Generate an SROM $\widetilde{Z}\,|\,\Gamma_i$ of $Z\,|\,\Gamma_i$ given by $\{\widetilde{z}_k\}_{k=1}^{m_i}$ and evaluate $U(\widetilde{z}_k)$.
6: Update the set $\mathcal{A} \leftarrow \mathcal{A} \cup \{\widetilde{z}_k\}_{k=1}^{m_i}$.
7: **end for**
8: **end while**

The size m_i of the SROM $\widetilde{Z}$ conditioned on Γ_i can be obtained by solving a constrained least squares problem. More specifically, the values of m_i can be chosen such that the proportion of samples $\widetilde{z}_k$ in Γ_i for which the response has been computed is equal to $P(\Gamma_i)$, the probability of the *partition* Γ_i. Let N be the number of partitions Γ_i in the current iteration and let $n_i := \#\{\widetilde{z}_j | \widetilde{z}_j \in \Gamma_i \cap \mathcal{A}\}$. The condition can then be mathematically expressed as $\dfrac{n_i + m_i}{\sum_{j=1}^{N} n_i + \sum_{j=1}^{N} m_i} = P(\Gamma_i)$ for $i = 1, \ldots, N$ which is a linear system with nonnegative constraints for m_i. A cheap estimate of $P(\Gamma_i)$ can be obtained as before.

For a practical implementation of the adaptive algorithm with local sampling, we have imposed a cap on the maximum number of response and gradient calculations for a given computational budget. This is because as the adaptive construction progresses, the partitions decrease in size and it becomes counterintuitive to compute more response evaluations to approximate the L^p error. By imposing a limit on the response and gradient calculations, the local sampling scheme then encounters the same problem as in the global sampling scheme wherein the L^p error may not be approximated for a particular *partition* for some iteration. Consequently, we address the issue similarly using contingency response evaluations. The stopping criterion under the local sampling approach is now similar as the global sampling approach.

Finally, using the L^p error as the refinement criterion has implications on the performance of the global and local sampling schemes we have proposed. The local sampling scheme allows one to obtain a surrogate at lesser expense since not all of the sample responses that are used to approximate the L^p error are generated during the initialization step. However, for some response functions and for some choices of the probability law for Z, the global sampling scheme can be favorable in that the approximate L^p errors of the partitions are more accurate at the early stages of the adaptive construction when the sizes of the partitions are large.

Remark. If a posteriori error estimates of $U(Z) - \widetilde{U}(Z)$ conditioned on Z are available, as in [5], the L^p error of the a posteriori error estimate can be used instead for the refinement measure. The L^p error gives us a metric over the probability space rather than at a single sample of Z only.

We now elaborate on some computational aspects and implementation issues of the adaptive algorithm.

3.1.3 Choice of parameters

The initialization step requires that values for p, ϵ need to be specified. Suppose that we are interested in performing a forward sensitivity analysis on the quantity of interest $q(Z) := \ell(U(x,t,Z))$ where $U(x,t,Z)$ is a solution to a PDE. If the objective is to examine the k-th order moments of $q(Z)$ for $k = 1, \ldots, n$, then it would be natural to set $p = n$ as convergence in L^p also guarantees convergence in L^q for $q \leq p$.

On the other hand, the choice of ϵ can be inferred from the following inequality. Let $U^h(x,t,Z)$ be the discretized solution of $U(x,t,Z)$ in the time and physical space and let $\widetilde{U}^h(x,t,Z)$ be the SROM-based surrogate based on $U^h(x,t,Z)$. If $q^h(Z) \coloneqq \ell(U^h(x,t,Z))$ and $\widetilde{q}^h(Z) \coloneqq \ell(\widetilde{U}^h(x,t,Z))$,

$$\|q(Z) - \widetilde{q}^h(Z)\| \leq \|q(Z) - q^h(Z)\| + \|q^h(Z) - \widetilde{q}^h(Z)\|.$$

The second term in the inequality right hand side is the error obtained by approximating the discretized solution by a surrogate and is precisely the stopping criterion of the adaptive method above. The first term however is due to the error of using a discretized solution in place of the analytic solution and bounds for the physical discretization error are usually available using a posteriori error estimates. Hence, ϵ must be chosen such that it is of the same order of magnitude as the error due to the physical discretization.

3.1.4 Selection of new node and choice of partitioning strategy

We now address how to select the new node $\widetilde{z}_j$ as outlined in line 7 of Algorithm 1 (and the corresponding line in Algorithm 2) and how to carry out the partitioning of Γ_k that is necessary for both algorithms.

Suppose that the partition with the largest approximate L^p error has been identified. Two schemes on the selection of the new node have been investigated, namely the surplus-based scheme and surplus-pdf-based scheme. Under the former scheme, we select the node $\widetilde{z}_j \in \Gamma_k \cap \mathcal{A}$ for which the surplus $|U(\widetilde{z}_j) - \widetilde{U}_m(\widetilde{z}_j)|$ is the largest. The latter scheme, however, selects the node for which $|U(\widetilde{z}_j) - \widetilde{U}_m(\widetilde{z}_j)|^p f(\widetilde{z}_j)$ is maximized and this quantity is precisely the integrand of the $L^p(\Omega)$ error assuming that Z has joint pdf $f(z)$. The surplus-based scheme has the advantage in that the joint pdf of Z need not be known explicitly. Despite this convenience, it may result in computing gradients in regions of low probability especially in the initial iterations, when the partitions are relatively large. This in turn can adversely affect the L^p error of the surrogate in the subsequent iterations.

We have also investigated two types of strategies for refining the partitions Γ_k which are the neighbor-based refinement and cell-based refinement. Using the notation in Algorithms 1 and 2, Figure 9(a) shows the original partition where the green asterisks denote $\widetilde{z}_k$ for which $\nabla U(\widetilde{z}_k)$ has been computed and the black asterisk is $\widetilde{z}_j$ which resides in the partition Γ_k. Under neighbor-based refinement, a Voronoi tessellation of Γ with $\widetilde{z}_j$ as an additional Voronoi center is performed to obtain the top right plot. Essentially, only neighboring Voronoi cells of Γ_k are affected by this partitioning strategy as the number of partitions increases even though this is not evident in Figure 9(b). To visualize this, one can consider an example in 2 stochastic dimensions with a unit square domain where the Voronoi cells are squares of identical sizes. It is clear that refining the upper right most Voronoi cell will not affect the

Voronoi cells in the lower left region of the domain. In contrast, for cell-based refinement, only a Voronoi tessellation conditioned on Γ_k with the black and the green asterisk as centers is performed to carry out the refinement in order to obtain the bottom plot. Hence, the resulting partitions in this case are not Voronoi cells.

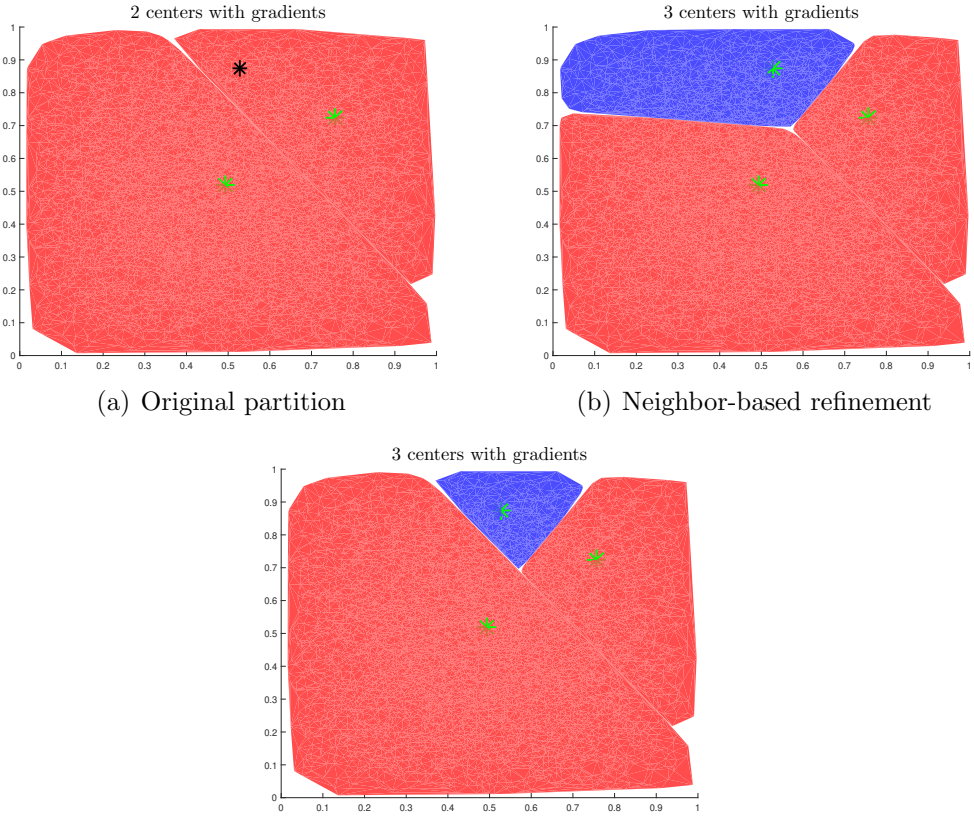

(a) Original partition (b) Neighbor-based refinement

(c) Cell-based refinement

Figure 9: Original partition (top left) and illustration of neighbor-based (top right) and cell-based (bottom) refinement.

From a computational perspective, it is clear that the cell-based refinement is less expensive and can be efficiently implemented using a tree data structure. As for the neighbor-based refinement, storing and updating the distance of each sample of z_i of Z (that we have generated to construct the SROM $\widetilde{Z}$) from its corresponding Voronoi center at every iteration will be convenient to speed up the sample clustering in the refinement process. However, we shall see later that the neighbor-based refinement strategy will be advantageous especially in high stochastic dimension.

An illustration of the adaptive method with global sampling using neighbor-based refinement is shown in Figure 10 where the surplus-based scheme was used to select the new node in a partition. As can be seen, the adaptive method prioritizes regions of Γ with high probability and with high variation in $U(z)$. In summary, the adaptive method and the

direct construction of the SROM-based surrogate as outlined in section 2.2.1 are similar in that both surrogates use the first-order Taylor interpolant in each Voronoi cell and both have the same number of gradient calculations. However, both methods differ in that the adaptive method has more evaluations of $U(z)$ in order to approximate the L^p error needed for the refinement criterion. In addition, the nodes $\widetilde{z}_j$ at which we compute $\nabla U(\widetilde{z}_j)$ in the adaptive method are located in regions of high probability in Z and large variation in $U(Z)$.

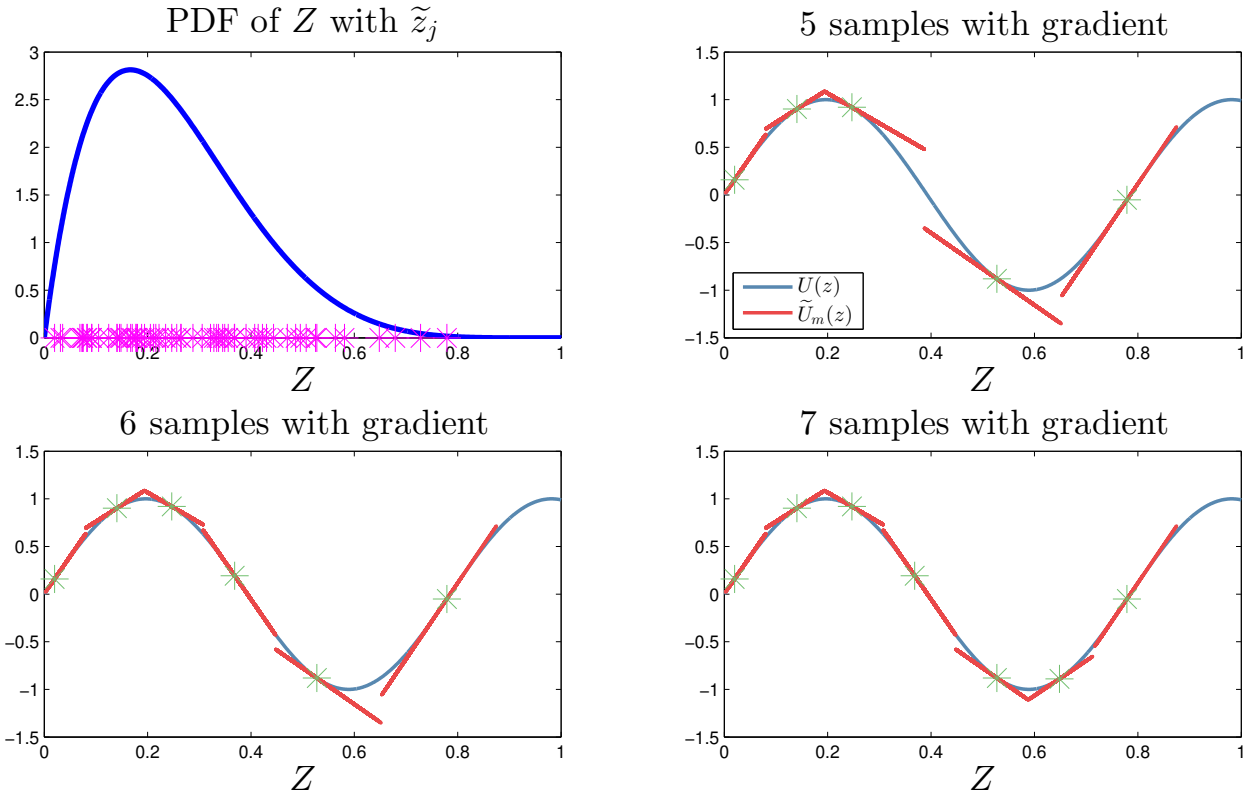

Figure 10: Illustration of the adaptive method with global sampling using neighbor-based refinement for $U(Z) = \sin(8Z)$, $Z \sim \text{beta}(2, 6)$, $p = 3$. The upper left subplot shows a plot of the pdf of Z (blue) and the SROM $\widetilde{Z}$ generated during the initialization step (magenta asterisks). For the remaining subplots, the blue curve is $U(z)$, the red curve is $\widetilde{U}_m(z)$, and the green asterisks are the nodes where the gradient is available.

3.1.5 Post-termination of the adaptive method

Upon termination of the adaptive algorithm, it is possible that $\mathcal{A} \neq \emptyset$, i.e. there are nodes $\widetilde{z}_j$ for which $U(\widetilde{z}_j)$ is known but for which there are no computational units left to obtain $\nabla U(\widetilde{z}_j)$. Although these nodes have served their purpose in providing validation about the SROM-based surrogate, it is natural to ask if this information can be used to improve the

quality of $\widetilde{U}_m(Z)$. The challenge is that we need to use an interpolant for these nodes where the gradient information is unavailable and that this interpolant must be compatible with the 1st-order Taylor interpolant on each partition.

We therefore propose the construction of a hybrid surrogate model as follows.

Algorithm 3 Hybrid SROM-based surrogate model

1: **for** each partition Γ_k **do**
2: Find $\{\widetilde{z}_j\} \in \mathcal{A} \cap \Gamma_k$ such that $|U(\widetilde{z}_j) - \widetilde{U}(\widetilde{z}_j)|^p > \epsilon$ or $|U(\widetilde{z}_j) - \widetilde{U}(\widetilde{z}_j)|^p f(\widetilde{z}_j) > \epsilon$.
3: Perform a cell-based refinement on Γ_k using $\{\widetilde{z}_j\}$ found in the previous step and the existing center of the partition Γ_k.
4: Select the nodes $\widetilde{z}_j$ from the first step such that its resulting partition within Γ_k is inside $\text{Conv}(\{\widetilde{z}_j\}_{j=1}^M)$, the set from the initialization step of the adaptive method.
5: Use Sibson's interpolant on these partitions corresponding to the nodes $\widetilde{z}_j$ found in the previous step and retain the Taylor-based interpolant on the complement of these partitions within Γ_k.
6: **end for**

This construction ensures that the hybrid surrogate is exact for linear responses and that it preserves the convergence properties of the 1st-order Taylor interpolant used in the adaptive method. An example of the hybrid surrogate is illustrated below for $U(Z) = \arctan(50 \cdot \|Z - 0.75\|^2)$, $Z \in \mathbb{R}^2$ where $Z_i \sim \text{beta}(2, 2)$ i.i.d, $p = 2$ in which Figure 11 exhibits the resulting partitions of Γ while Figure 12 illustrates the surrogates. In Figure 11, the left subplot refers to the resulting partition once the adaptive method with global sampling using neighbor-based refinement terminates while the right subplot exhibits the partitioning of the hybrid surrogate with Sibson's interpolant employed in the green partitions. The interpolant in the red partitions is unchanged and is constructed using the 1st-order Taylor interpolant. All the asterisks in the subplots correspond to the centers of their corresponding partitions and the contour lines of $U(z)$ are also displayed in all subplots.

As Figure 12 demonstrates, Sibson's interpolant is able to smooth out the overshoot caused by using the 1st-order Taylor interpolant as can be observed by comparing the region $(z_1, z_2) \in [0.4, 1] \times [0, 0.5]$ in both subplots. This is because for a fixed partition, the gradient of Sibson's interpolant is not constant unlike that of the 1st-order Taylor interpolant, making it more suitable for response surfaces with large curvature. This is quantitatively manifested in the L^p error plots in Figure 13 where the horizontal axis represents the iteration number or equivalently, the number of nodes with gradient calculated. The L^p error of the hybrid surrogate is plotted at the last iteration and the error is lower than that of the previous iteration corresponding to the termination of the adaptive algorithm. We acknowledge that we cannot prove that the hybrid surrogate reduces the L^p error of the surrogate obtained from the adaptive method. However, we have the consolation that Sibson's interpolant guarantees that the surrogate does not go beyond the minimum and maximum values of the response according to Property 9.

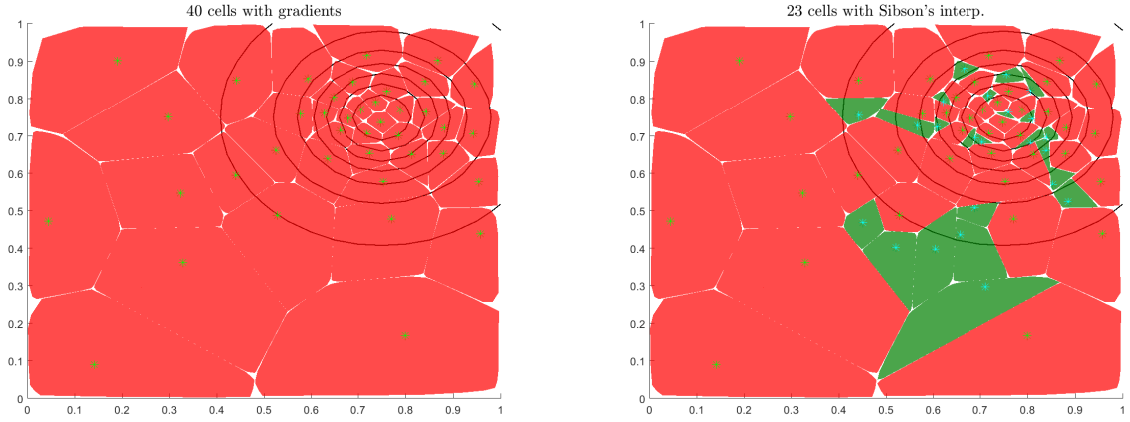

(a) Original partition (b) Hybrid partition

Figure 11: Partition from the adaptive method (left) and for the hybrid surrogate (right).

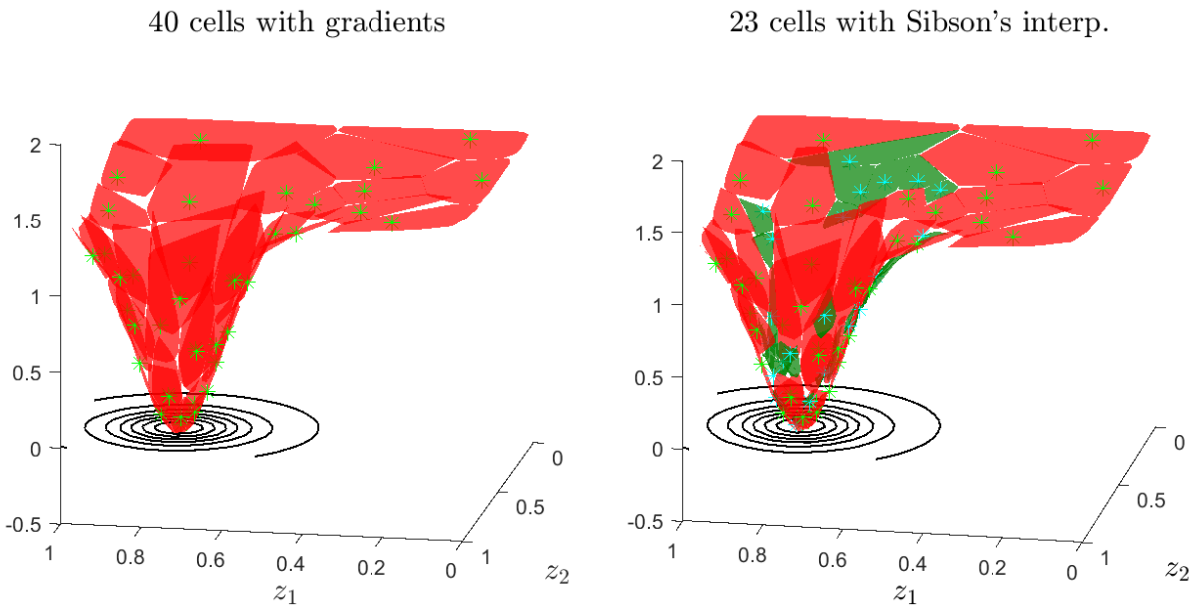

(a) Original surrogate (b) Hybrid surrogate

Figure 12: SROM-based surrogate obtained from the adaptive method (left) and hybrid SROM-based surrogate (right).

3.2 Analysis of the algorithm

As can be observed from Figure 13, the L^p error of the SROM-based surrogate obtained from the adaptive method is not monotonically decreasing as a function of the number of iterations. The number of iterations is the same as the number of nodes with gradient, hence,

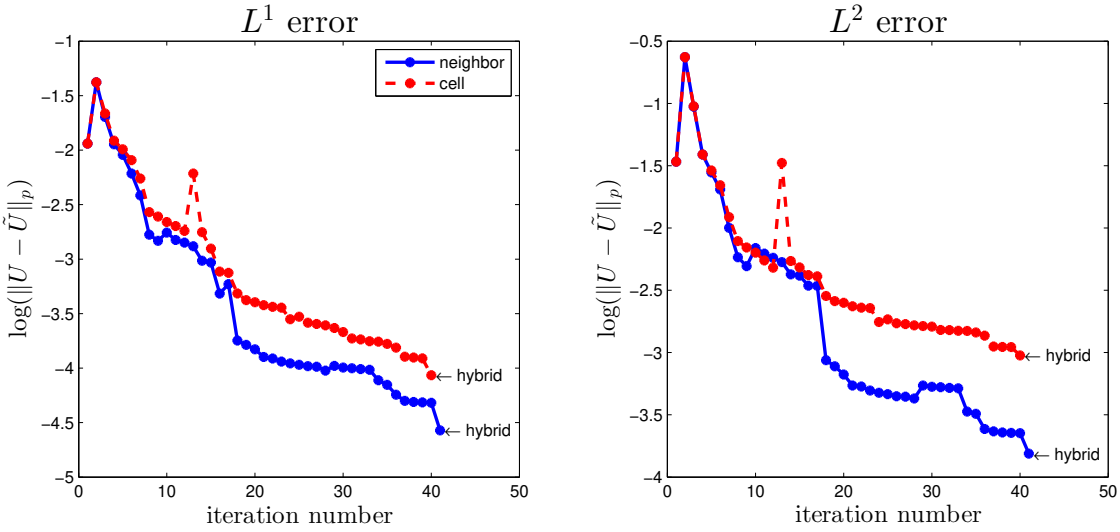

Figure 13: L^p error of the SROM-based surrogate and the hybrid surrogate for $p = 1$ (left) and $p = 2$ (right) as a function of the number of nodes with ∇U computed. Results for neighbor-based refinement is shown in red while that for cell-based is shown in blue.

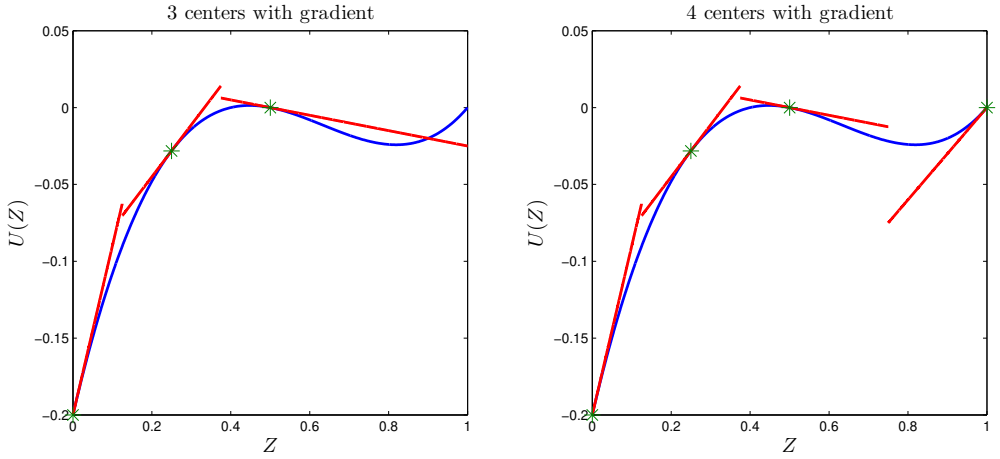

Figure 14: Illustration as to why the L^p is not monotonically decreasing.

this result appears counterintuitive given that more computations have been expended. We now provide a simple analytic example to demonstrate why such situation occurs.

Example 3. Let $U(Z) = (Z - 0.4)(Z - 0.5)(Z - 0.1)$ where $Z \sim \text{Unif}(0, 1)$ and $p = 1$. In Figure 14, an illustration of how the adaptive method using neighbor-based refinement transitions from the 3rd iteration (left) to the 4th (right) is presented. The nodes with gradient in iteration 3 are $\{\widetilde{z}_j\}_{j=1}^3 = \{0, 0.25, 0.5\}$ while $\widetilde{z}_4 = 1$ is the additional node with gradient in iteration 4.

In order to see why the L^1 norm is not monotonically decreasing, it suffices to compare $\int_{0.75}^{1} |U(z) - \widetilde{U}_3(z)|\, dz$ and $\int_{0.75}^{1} |U(z) - \widetilde{U}_4(z)|\, dz$ because $z \in [0.75, 1]$ is the portion of the domain for which the interpolant has changed. On the interval $[0.75, 1]$, $\widetilde{U}_3(z) = U(0.5) + U'(0.5)(z - 0.5)$ whereas $\widetilde{U}_4(z) = U(1) + U'(1)(z - 1)$. In addition, using Taylor's remainder theorem, we know that

$$U(z) - \widetilde{U}_j(\widetilde{z}_j) = \frac{1}{2} U''(\xi(z))(z - \widetilde{z}_j)^2$$

for some ξ which is a function of $z \in [0.75, 1]$. Hence, elementary calculations show that $|U(z) - \widetilde{U}_3(z)| = |(2z - 1.8)(z - 0.5)^2|$ while $|U(z) - \widetilde{U}_4(z)| = |(2z + 0.2)(z - 1)^2|$. Even though $|z - 1|^2 < |z - 0.5|^2$ for $z \in [0.75, 1]$, this is offset by $|2z - 1.8| < |2z + 0.2|$ which implies that $\|U(z) - \widetilde{U}_4(z)\|_{L^1} > \|U(z) - \widetilde{U}_3(z)\|_{L^1}$.

The example above leads to the following proposition which provides a sufficient condition to ensure that the L^p error is monotonically decreasing.

Proposition 10. *Suppose that the Hessian of $U(z)$ is constant and that all its eigenvalues are all equal. It then follows that $\|U - \widetilde{U}_{m+1}\|_{L^p} \leq \|U - \widetilde{U}_m\|_{L^p}$, i.e. the L^p error is monotonically decreasing.*

Proof. Under the stated assumptions, the discrepancy between the response and the surrogate on a partition Γ_k as shown in (2.5) can be simplified into the following. As before, $y \coloneqq Q \cdot (z - \widetilde{z}_k)$ where Q is an orthonormal matrix in the eigendecomposition of the Hessian and $\widetilde{z}_k$ is the center of the partition. We have that

$$|U(z) - \widetilde{U}_m(z)| = \frac{1}{2} \left| \sum_{i=1}^{d} \lambda_i y_i^2 \right| = \frac{|\lambda|}{2} \|Q \cdot (z - \widetilde{z}_k)\|^2 = \frac{|\lambda|}{2} \|z - \widetilde{z}_k\|^2$$

where λ is the constant eigenvalue.

Hence the L^p error of the surrogate on a partition is simplified as

$$\|U(z) - \widetilde{U}_m(z)\|_{L^p(\Gamma_k)}^p = \left(\frac{|\lambda|}{2} \right)^p \int_{\Gamma_k} \|z - \widetilde{z}_k\|^{2p}\, dF(z).$$

As shown in the previous example, it only suffices to examine subsets of the partition Γ_k for which the interpolant over the subset changed from one iteration to the next. Let Γ' be such

a subset such that $\Gamma' \subset \Gamma_j$ in iteration m but $\Gamma' \subset \Gamma_k$ in iteration $m+1$. Denote the center of the partition Γ_i by $\widetilde{z}_i$.

It therefore follows that

$$\begin{aligned}\|U(z) - \widetilde{U}_{m+1}(z)\|^p_{L^p(\Gamma')} &= \left(\frac{|\lambda|}{2}\right)^p \int_{\Gamma'} \|z - \widetilde{z}_k\|^{2p} \, dF(z) \\ &\leq \left(\frac{|\lambda|}{2}\right)^p \int_{\Gamma'} \|z - \widetilde{z}_j\|^{2p} \, dF(z) \\ &= \|U(z) - \widetilde{U}_m(z)\|^p_{L^p(\Gamma')}\end{aligned}$$

where the inequality holds by definition of Voronoi tessellation.

□

In 1-dimension, the sufficient conditions are tantamount to the response having constant concavity. We also remark that the statement above holds regardless of the partitioning strategy chosen for the adaptive method. Even though the L^p error is not monotonically decreasing in general, we expect that the L^p error will have a decreasing trend due to the convergence properties of the SROM-based surrogate.

4 Numerical results

We now demonstrate the benefits of an adaptive construction of the SROM-based surrogate over the direct method through a variety of numerical examples in different stochastic dimensions. Since we expect that the adaptive method will be more computationally expensive, we have ensured that the number of gradient calculations for both methods is identical for the comparison to be fair. In addition, as mentioned above, the SROM-based surrogate obtained is not unique but we do not attempt to generate multiple instances of the SROM-based surrogate in the following results.

In Example 4, a 2-d example is considered where the components of the random vector $Z = (Z_1, Z_2)$ are uniform i.i.d. in order to isolate the effect of the variation of $U(Z)$. We then compare the resulting partitioning of $\Gamma = Z(\Omega)$ obtained using direct and adaptive constructions. We continue this example with Example 5 wherein we compare the resulting partitioning of Γ for the global and local sampling scheme. Examples 6 and 7 are high-dimensional examples which investigate the performance of the direct and adaptive method with global sampling in cases where the region of high probability and the region of large variation in $U(Z)$ exactly coincide or slightly intersect. On the other hand, Examples 8 and 9 consider an application to solving a partial differential equation with random parameters

in high dimensions. Lastly, a comparison between the performance of the adaptive algorithm with global and local sampling is carried out in Example 10.

In the following two examples, we visually compare the resulting partitioning of Γ for various methods of constructing the SROM-based surrogate. For instance, in Example 4, we aim to show that the adaptive method is able to track regions of high variation unlike the direct construction. In Example 5, we show how the resulting SROM nodes under local sampling can be more spread out than that produced by global sampling.

Example 4. Consider $U(Z) = \arctan(50 \cdot \|Z - 0.75\|^2)$, $Z \in \mathbb{R}^2$ where $Z_i \sim \text{Unif}(0,1)$ i.i.d, $p = 3$. An illustration of the response is shown in Figure 12 above. Figure 15 demonstrates the resulting 40 partitions of Γ using direct construction (left), adaptive method with neighbor-based refinement (middle), and adaptive method with cell-based refinement (right), with the contour lines of $U(z)$. Both adaptive algorithms were carried out with global sampling. With only 84 additional evaluations of $U(z)$ for both refinement strategies, the adaptive method targets regions where the variation of $U(z)$ is large (since the distribution of Z in this case is uniform).

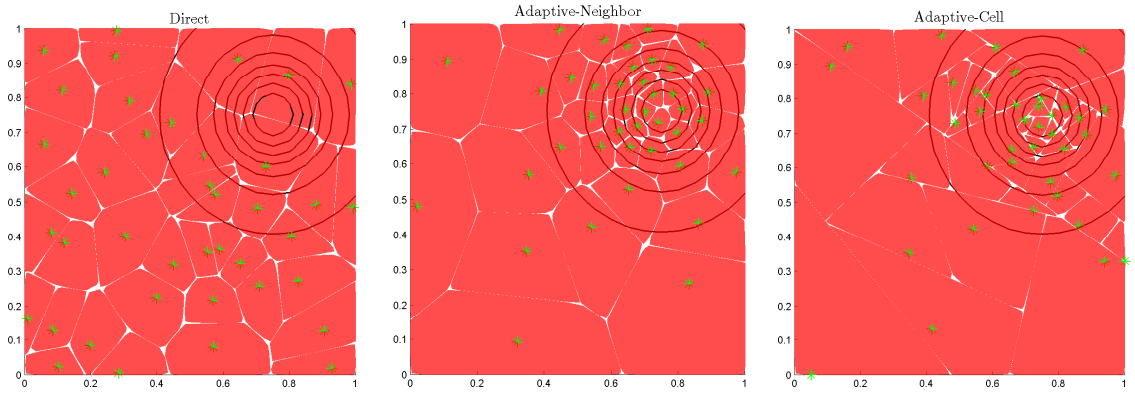

Figure 15: Comparison of the partitioning of Γ for 3 types of construction of SROM-based surrogates: Direct, Adaptive-Neighbor, Adaptive-Cell (Left, Middle, and Right panels).

Example 5. Let $U(Z)$ be as in Example 4 with the same probability law for Z. Figure 16 demonstrates the resulting 30 partitions of Γ for the adaptive method using neighbor-based refinement with global sampling (left) and local sampling (right). All of the asterisks in both figures represent the nodes $\widetilde{z}_k$ for which $U(z)$ has been evaluated. In addition, the green asterisks represent the nodes for which $\nabla U(z)$ has been computed. A comparison of the subplots shows how the nodes acquired through local sampling are more spread out compared to that of global sampling because some of these nodes were selected as the adaptive construction progressed. Quantitatively, the marginal variance along Z_1 for the nodes obtained from global sampling is 0.0593 while the marginal variance along Z_2 is 0.0765. For comparison, the marginal variances for the nodes under local sampling are 0.0779 and 0.0829 along Z_1 and Z_2, respectively.

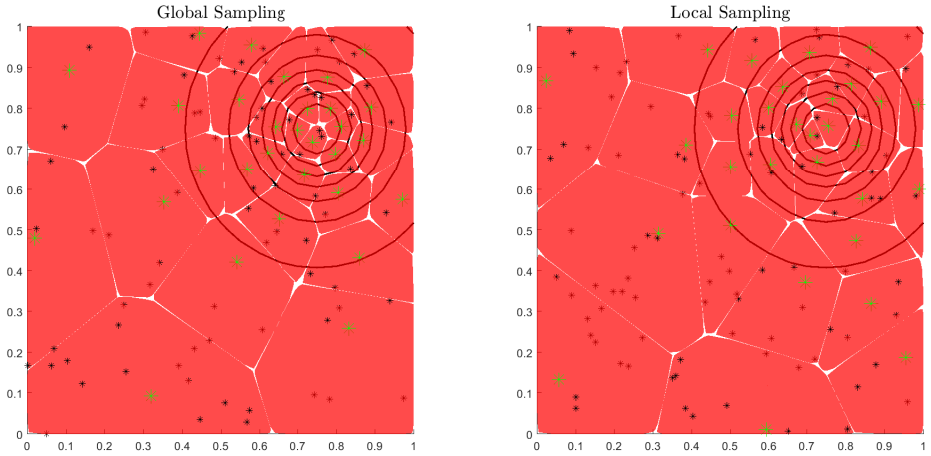

Figure 16: Comparison of the partitioning of Γ for 2 types of construction of SROM-based surrogates: Global sampling and Local sampling (Left and Right panels).

We now examine two related scenarios in high stochastic dimension. We first consider the case when the region where Z has high probability coincides with the region where $U(Z)$ has high variation as the next example illustrates. In this scenario, we expect the surrogate resulting from the direct construction to perform well as the need to locate regions of high variation is reduced. Subsequently, we will investigate the case when the region of high probability of Z is mostly concentrated on regions where $U(Z)$ has low variation. For both of these examples, global sampling was used for the adaptive algorithm.

Example 6. Let $U(Z) = \arctan(50 \cdot \|Z - 0.5\|^2)$, $Z \in \mathbb{R}^6$ where $Z_i \sim F^{-1}(\Phi(Y_i))$, $Y_i \sim N(0,1)$, $\text{cov}(Y_i, Y_j) = 0.2$ for $i \neq j$, F is the CDF of beta$(30, 30)$, and $p = 4$. Figure 17 (left) shows the contour plot of the 2-d cross section of $U(z)$ for $z_i = 0.5, i > 2$ with samples of (Z_1, Z_2).

The L^p error of the SROM-based surrogate obtained under direct construction (dashed green), adaptive method using surplus-pdf-based scheme (thick solid blue), and adaptive method using surplus-based scheme (thin solid red) is shown in Figure 18 for $p = 1, \ldots, 4$ as a function of the number of nodes with gradient calculated. Neighbor-based refinement was used for both adaptive methods. The L^p error under the direct construction was obtained using nodes with gradient calculations (equivalently, the number of partitions) in increments of 5. In addition, these sets of nodes do not form a refining sequence unlike in the case for the adaptive methods, i.e. the SROM nodes in the previous iteration is not a subset of the nodes in the succeeding iteration. With 235 extra evaluations of $U(z)$ for the adaptive surplus-based method and 237 extra evaluations for the surplus-pdf based method in the last iteration, we see the benefits of using an adaptive approach compared to a direct construction. In this example, there is no significant difference in the performance of the two schemes for selecting a new node.

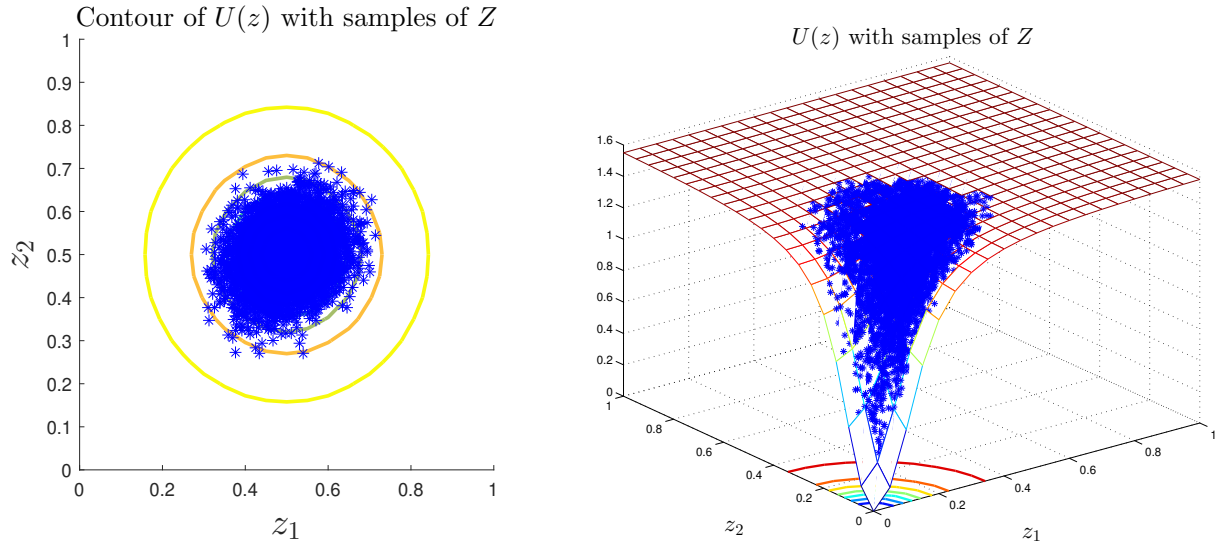

(a) Example 6 (b) Example 7

Figure 17: Illustration of the response and distribution of the random vector for Example 6 (left) and for Example 7 (right).

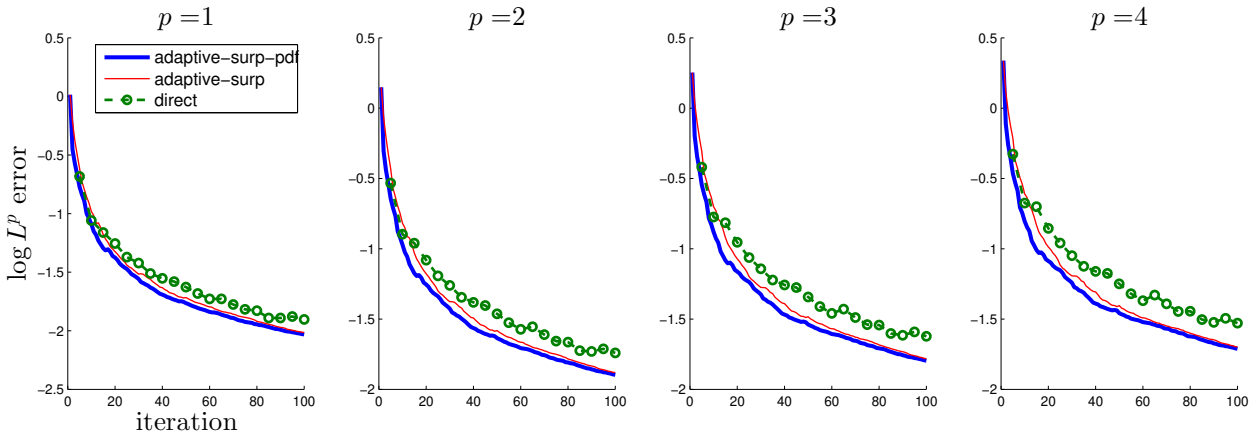

Figure 18: L^p error of the surrogate for the response in Example 6 obtained using 3 types of construction: Direct, Adaptive-surplus-pdf, Adaptive-surplus (Dashed, Thick Solid, Thin Solid).

Example 7. Set $U(Z) = \arctan(50 \cdot \|Z\|^2)$, $Z \in \mathbb{R}^{20}$ where $Z_i \sim F^{-1}(\Phi(Y_i))$, $Y_i \sim N(0,1)$, $\text{cov}(Y_i, Y_j) = 0.4$ for $i \neq j$, F is the CDF of beta$(6, 20)$, and $p = 4$. An illustration of the 2-d cross-sectional response $U(z)$ for $z_i = 0, i > 2$ and samples of (Z_1, Z_2) is shown in Figure 17 (right).

The L^p errors of the surrogate under the 3 different types of construction are presented

in Figure 19 with the surplus-based scheme requiring 513 additional evaluations of $U(z)$ while the surplus-pdf-based scheme only requiring 503 more for the last iteration. Both adaptive methods outperform the direct construction in general, especially as p increases. For example, at the 80th iteration, there is a 62% decrease in the L^4 error if the adaptive surplus-pdf-based scheme is used instead of the direct construction.

The fluctuations in the higher L^p error under the direct construction result from the amplification of the error between the response and the surrogate that is not explicit in the L^1 error. Furthermore, this demonstrates the challenges associated with approximating response surfaces with tangent hyperplanes as was discussed in Section 3.2 for the case when the SROM nodes form a refining sequence. Even in the case where the SROM nodes of the preceeding iteration is not a subset of the succeeding iteration, as in the direct construction here, it is still possible that the L^p norm will not be monotonically decreasing. This is because of the lack of curvature of the tangent hyperplane approximation.

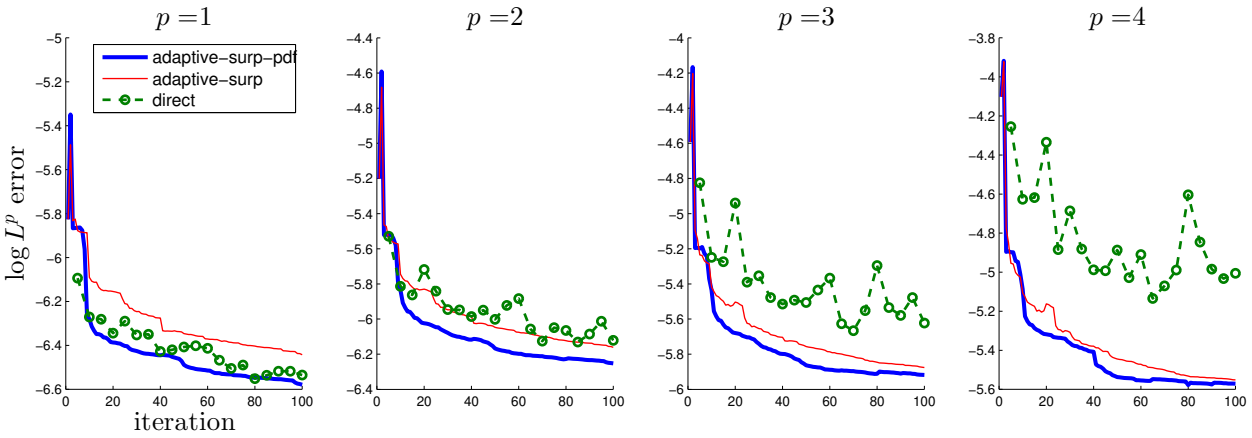

Figure 19: L^p error of the surrogate for the response in Example 7 obtained using 3 types of construction: Direct, Adaptive-surplus-pdf, Adaptive-surplus (Dashed, Thick Solid, Thin Solid).

In the examples above, we did not display the L^p error of the adaptive method using cell-based refinement as it has the tendency to underperform compared to the direct construction despite the latter not taking into account information about regions where $U(z)$ has high variation. This is because only 1 partition changes from one iteration to the next under cell-based refinement which makes it prone to slow convergence especially if the response is very smooth. For comparison, the number of partitions that change between consecutive iterations can be large under neighbor-based refinement especially in high dimensions. To visualize this, the simplest Voronoi cell is a hyperrectangle which has $2d$ faces in d dimensions. We quantify these observations through Figure 20 where a comparison is made for the L^4 error of the surrogate obtained using neighbor-based refinement (thick blue) and cell-based refinement (thin red) for Examples 6 and 7. Both schemes for selecting a new node have

been considered and it is evident that the disparity can be considerable for both types of refinement. As a consequence, we will only focus on the neighbor-based refinement in the examples that follow.

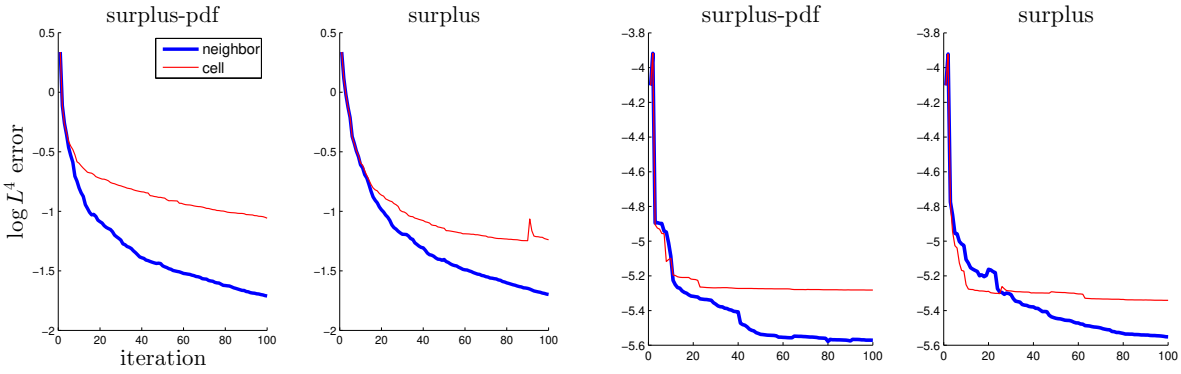

(a) Example 6 (b) Example 7

Figure 20: Comparison of the L^4 error using the neighbor-based refinement (thick blue) and cell-based refinement (thin red) using both schemes of selecting a new node for Example 6 (left 2 subplots) and for Example 7 (right 2 subplots).

We then present examples in which the response is a solution to a stochastic partial differential equation as given in Examples 8 and 9. As before, we employ the global sampling scheme for the adaptive algorithm.

Example 8. Consider the PDE $\nabla \cdot (A(x,Z)\nabla U(x,Z)) = 0$, $x \in D$ with $D = (0,l_1) \times (0,l_2)$ and the boundary conditions $U(0,x_2) = 0, U(l_1,x_2) = 1$ for $x_2 \in (0,l_2)$ and $U_{x_2}(x_1,0) = U_{x_2}(x_1,l_2) = 0$ for $x_1 \in (0,l_1)$. $Z = (Z_1,\ldots,Z_d)$ is a random vector defined on the probability space $(\Omega,\mathcal{F},P)$. Physically, $U(x,Z)$ represents the electric potential on a specimen with $A(x,Z)$ being the conductivity field. A quantity of interest is the apparent conductivity given by

$$q(Z) = \frac{1}{l_2}\int_D A(x,Z)\frac{\partial U(x,Z)}{\partial x_1}\,dx$$

which is a random variable defined on the same probability space as Z. Our objective is to construct an SROM-based surrogate for $q(Z)$.

Suppose that $Z \in \mathbb{R}^4$ where $Z_i \sim \Phi(Y_i)$, $Y_i \sim N(0,1)$, $\mathrm{cov}(Y_i,Y_j) = 0.9$ for $i \neq j$, and $p = 4$, and that our conductivity field is given by

$$A(x,Z) = 4 + \sum_{i=1}^{2}(\sin(2\pi Z_{2i-1})\cos(2ix_1+2ix_2) + \cos(2\pi Z_{2i})\sin(2ix_1+2ix_2)).$$

The left subplot of Figure 21 shows samples of (Z_1, Z_2) whereas the right subplot shows the response surface of $q(Z_1, Z_2, 0.5, 0.5)$ using $l_1 = l_2 = 1$.

We remark that the construction of an SROM-based surrogate for $q(Z)$ requires calculations of $\dfrac{\partial U}{\partial Z_i}$ which can be obtained by differentiating both sides of the PDE with respect to Z_i and solving the resulting PDE.

Figure 22 shows the L^p errors of the SROM-based surrogate constructed in 3 different ways as before. The surplus-based scheme required 173 additional evaluations of $U(z)$ while the surplus-pdf-based scheme only required 150 more for the last iteration. The results show that the gap between the L^p errors of the adaptive and direct construction increases as p increases. For example, at the 55th iteration, there is a 34% decrease in the L^4 error if the adaptive surplus-pdf-based scheme is used instead of the direct construction.

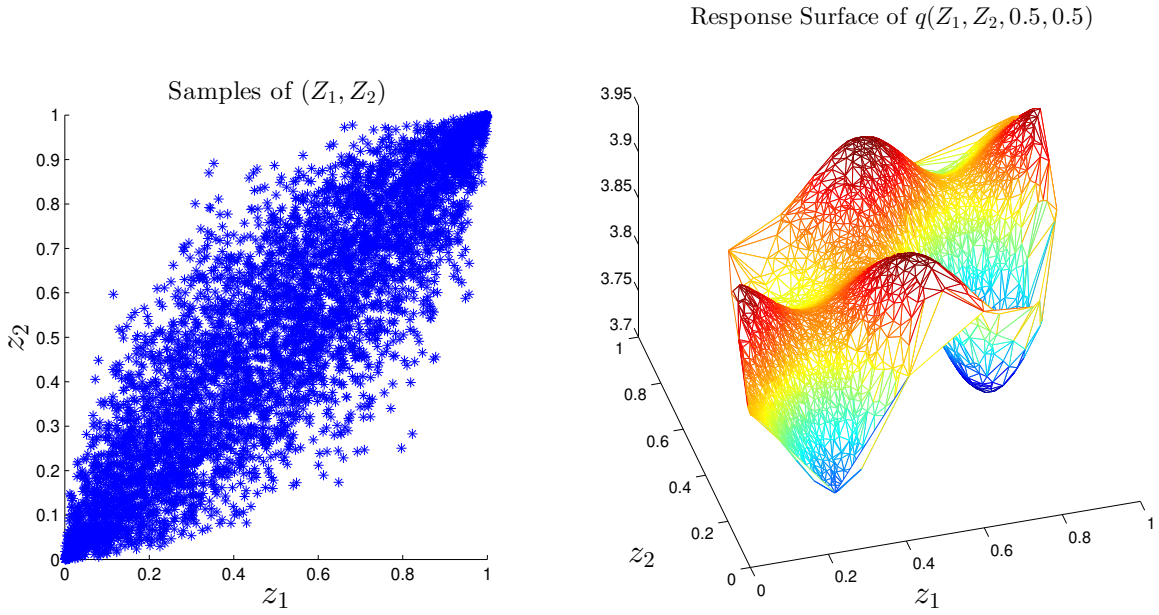

Figure 21: Illustration of the response and distribution of the random vector for Example 8.

Example 9. Consider the PDE in Example 8 using the same boundary conditions for $l_1 = l_2 = 0.2$. We construct the conductivity field as $A(x, Z) = F^{-1} \circ \Phi(G(x, Z))$ where F is the CDF of a Beta distribution with support $[1, 8]$ and shape parameters 2 and 6. Φ represents the CDF of the standard normal random variable with G being a homogeneous Gaussian field with mean 0, variance 1 and covariance function $c(\xi) = E[G(x, Z)G(x + \xi, Z)] = \exp(-0.5 \cdot (\xi_1^2 + 2\rho\xi_1\xi_2 + \xi_2^2))$ where $\xi = (\xi_1, \xi_2) \in \mathbb{R}^2$ with $\rho = 1000$. G is parametrically represented by a truncated Karhunen-Loève expansion taking on the form $G(x, Z) = \sum_{i=1}^{20} \sqrt{\lambda_i}\phi_i(x)Z_i$ where λ_i and ϕ_i are the eigenvalues and eigenfunctions of $c(\xi)$, respectively, and $Z_i \sim N(0, 1)$, *i.i.d.* As before, we are still interested in constructing a surrogate for $q(Z)$ where $Z \in \mathbb{R}^{20}$.

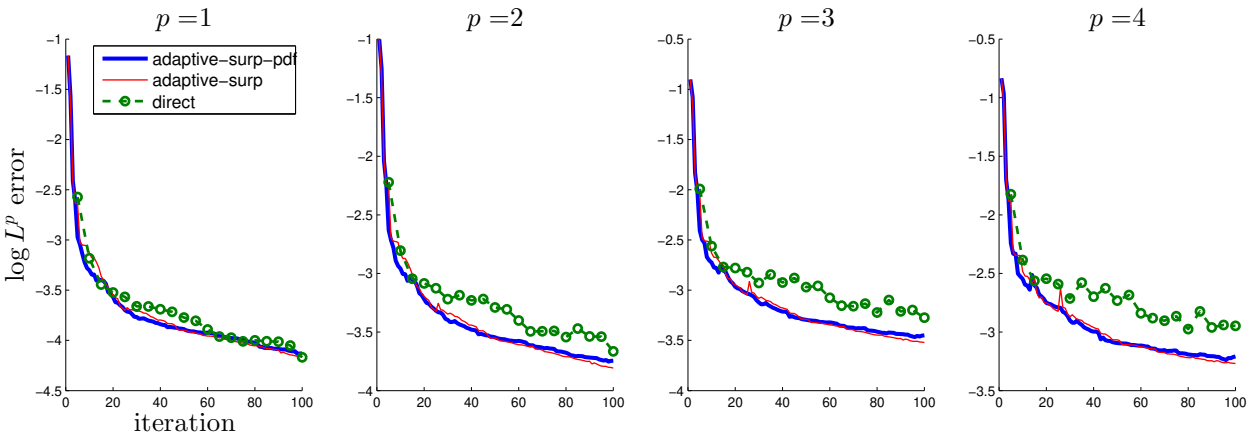

Figure 22: L^p error of the surrogate for the response in Example 8 obtained using 3 types of construction: Direct, Adaptive-surplus-pdf, Adaptive-surplus (Dashed, Thick Solid, Thin Solid).

We visualize the response surface in Figure 23 by constructing cross-sections of $q(Z)$ along coordinate axes Z_i and Z_j for which λ_i, λ_j are of similar magnitude. The values at the remaining coordinates are set to zero, the mean of Z_i. The plots indicate that there is a larger variation in the response outside the region of high probability. On the other hand, the cross-sections of $q(Z)$ along coordinate axes where $\lambda_i \gg \lambda_j$ appear to be linear, facilitating the approximation by tangent hyperplanes in these cases.

Figure 24 compares the L^p error of the surrogate under the 3 types of construction that we have considered in the previous examples. The direct construction of the SROM-based surrogate was performed in increments of 10 iterations where the SROM nodes do not form a refining sequence as before. The adaptive surplus-based scheme has 903 extra evaluations of $U(Z)$ compared to the direct construction for the last iteration while the adaptive surplus-pdf based scheme has 850 more than the direct construction. The advantage of the surplus-pdf-based scheme in selecting the new node is quantified in Figure 24. The performance of the surplus-based scheme is hampered by the fact that the regions of high variation are mostly outside the region of high probability. As a consequence, in the earlier iterations of the adaptive method when the partitions are still large, the surplus-based scheme computes gradients in regions of low probability. Because the curvature of the response in these regions tends to be larger, the resulting surpluses also tend to be larger, resulting in a surrogate that prioritizes this region in the earlier stages of the algorithm. This is because, in certain cases, a very large surplus will overshadow the impact of the factor $P(\Gamma_k)$ in the L^p error expression given by $E[|U(z) - \widetilde{U}_m(z)|^p \,|\, Z \in \Gamma_k] \cdot P(\Gamma_k)$ especially for large p which implies that partitions of smaller probability are refined in the initial stages of the adaptive construction. These factors combined affect the performance of the surrogate in the succeeding iterations, yielding a slower convergence in this example.

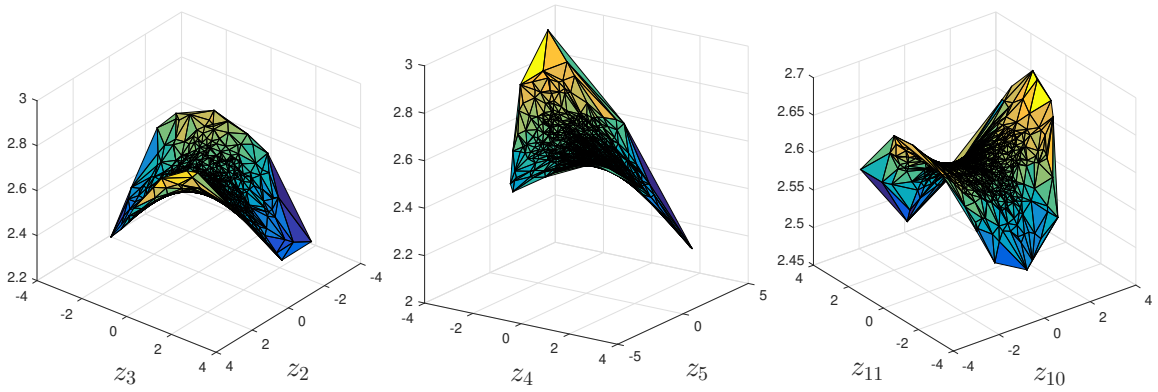

Figure 23: Cross-sections of the response $q(Z)$, $Z \in \mathbb{R}^{20}$, in Example 9.

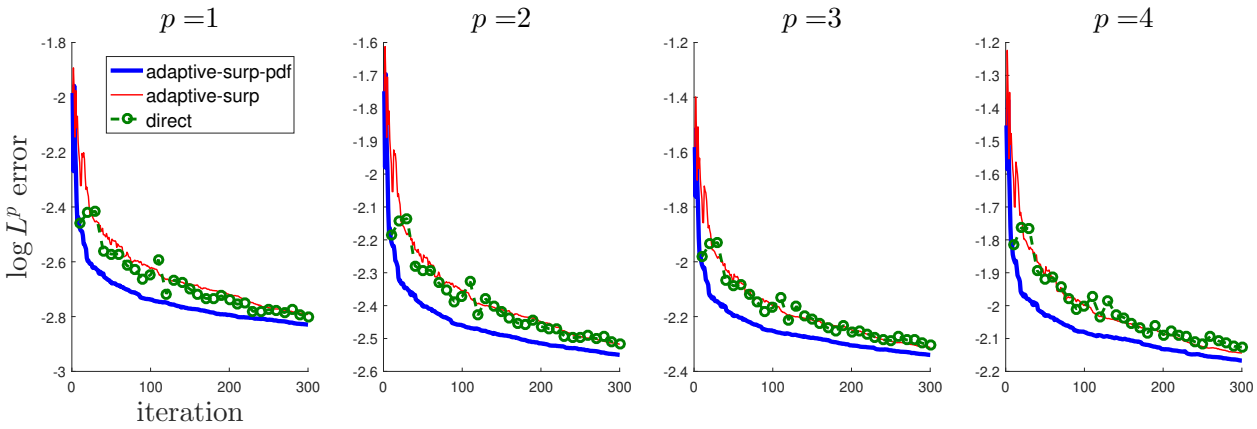

Figure 24: L^p error of the surrogate for the response in Example 7 obtained using 3 types of construction: Direct, Adaptive-surplus-pdf, Adaptive-surplus (Dashed, Thick Solid, Thin Solid).

We conclude with an example in which we are interested in comparing the performance of the two types of sampling schemes we have proposed, namely global and local sampling.

Example 10. We revisit the response $q(Z)$ in Example 8 with the corresponding probability law for Z. The adaptive construction was carried out using the surplus-pdf based scheme in selecting a new node.

Figure 25 shows the L^p error of the surrogates as a function of the computational budget. This is defined as the total number of response and gradient calculations for each surrogate at each iteration. We will assume that calculating a partial derivative is as costly as evaluating the response which implies that the computational budget in this example refers to the total

number of times a PDE was solved in each iteration. For this example, the global sampling scheme was initialized with 100 evaluations of the response $U(z)$ while the local sampling scheme was initialized with 30 evaluations of $U(z)$ with additional samples added as the iteration progressed. As can be seen, the local sampling scheme produces a surrogate with decent performance at a lesser expense at the beginning stages of the adaptive algorithm. However, in some cases such as this example, the global sampling scheme provides a more accurate surrogate in that the L^p errors are better approximated at the initial stages of the adaptive construction. In general, the performance between the two sampling schemes depends on the response function as well as the probability law of Z.

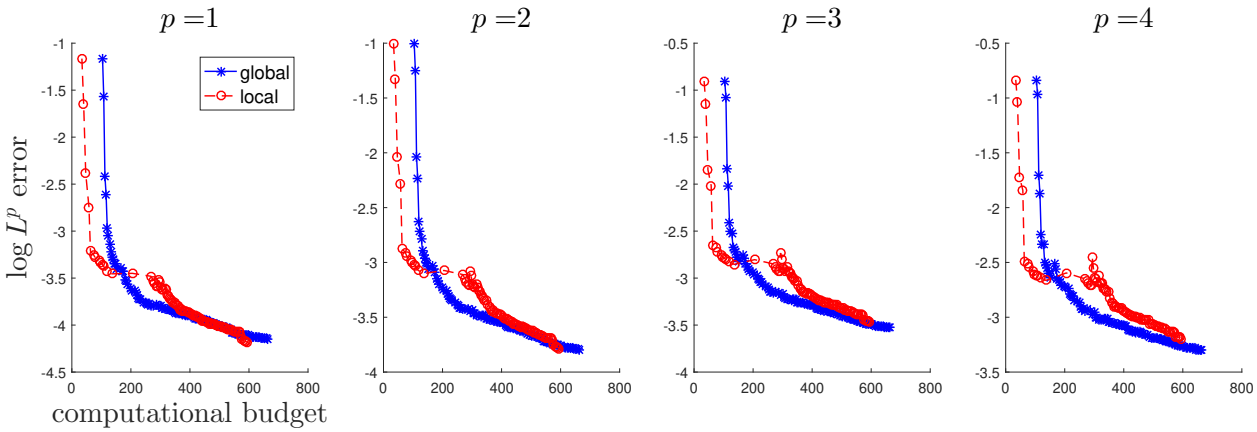

Figure 25: L^p error of the surrogate for the response in Example 7 obtained using 3 types of construction: Global sampling and Local sampling (Solid and Dashed).

5 Conclusion

We have surveyed novel and commonly used collocation-based surrogates for forward sensitivity analysis. We examined how the performance of sparse grid collocation can be adversely affected in high stochastic dimension as it does not take into account the probability law of the random input vector in the construction of the surrogate. To ameliorate this deficiency, we advocate the use of the SROM-based surrogate in which the interpolants employed are compatible with collocation nodes that are statistically representative of the input distribution. A comparison between zeroth-order Taylor, first-order Taylor, and the newly introduced Sibson's interpolant, was made in terms of their implementation and convergence properties.

Since the direct construction of the SROM-based surrogate does not take into account the regions where the response exhibits large variations, we proposed an adaptive construction which simultaneously targets regions with high probability and high variation. For this adaptive method, two approaches for selecting the new node in a partition, two approaches for

sequentially partitioning the probability space, and two sampling strategies were proposed and investigated, and their implementation and rate of convergence were compared. Subsequently, we analyzed how the adaptive method impacts the L^p error of the surrogate. Finally, numerical examples were furnished to demonstrate the benefits of the adaptive construction using various responses and probability distributions in different stochastic dimensions.

Acknowledgement

Wayne would like to thank Haoran Zhao and Christopher Earls for the helpful discussions. The authors are also grateful to the anonymous referees who provided useful feedback.

References

[1] Mircea Grigoriu, *Stochastic Systems: Uncertainty Quantification and Propagation*, Springer, 2012.

[2] Mircea Grigoriu, *Response Statistics for Random Heterogeneous Microstructures*, SIAM/ASA J. Uncertainty Quantification 2 (2014) 252-275.

[3] Jeroen A. S. Witteveen and Gianluca Iaccarino, *Simplex Stochastic Collocation with Random Sampling and Extrapolation for Nonhypercube Probability Spaces*, SIAM J. Sci. Comput. 34 (2) (2012) A814-A838.

[4] Jeroen A. S. Witteveen and Gianluca Iaccarino, *Refinement Criteria for Simplex Stochastic Collocation with Local Extremum Diminishing Robustness*, SIAM J. Sci. Comput. 34 (3) (2012) A1522-A1543.

[5] J.D. Jakeman and T. Wildey, *Enhancing Adaptive Sparse Grid Approximations and Improving Refinement Strategies Using Adjoint-based A Posteriori Error Estimates*, J. Comput. Phys. 280 (2015) 54-71.

[6] Xiang Ma and Nicholas Zabaras, *An Adaptive Hierarchical Sparse Grid Collocation Algorithm for the Solution of Stochastic Differential Equations*, J. Comput. Phys. 228 (2009) 3084-3113.

[7] Howard C. Elman and Christopher W. Miller, *Stochastic Collocation with Kernel Density Estimation*, Comput. Methods Appl. Mech. Engrg. 245-246 (2012) 36-46.

[8] Ph. Bressolette, M. Fogli, and C. Chauvière, *A Stochastic Collocation Method for Large Classes of Mechanical Problems with Uncertain Parameters*, Probabilist. Eng. Mech. 25 (2010) 255-270.

[9] Dongbin Xiu and Jan S. Hesthaven, *High-order Collocation Methods for Differential Equations with Random Inputs*, SIAM J. Sci. Comput. 27 (3) (2005) 1118-1139.

[10] F. Nobile, R. Tempone, and C. G. Webster, *A Sparse Grid Stochastic Collocation Method for Partial Differential Equations with Random Input Data*, SIAM J. Numer. Anal. 46 (5) (2008) 2309-2345.

[11] Jasmine Foo, Xiaoliang Wan, George Em Karniadakis, *The Multi-element Probabilistic Collocation Method (ME-PCM): Error Analysis and Applications*, J. Comput. Phys. 227 (2008) 9572-9595.

[12] R.V. Field Jr., M. Grigoriu, and J.M. Emery, *On the Efficacy of Stochastic Collocation, Stochastic Galerkin, and Stochastic Reduced Order Models for Solving Stochastic Problems*, Probabilist. Eng. Mech. 41 (2015) 60-72.

[13] R. V. Field and M. Grigoriu, *Convergence Properties of Polynomial Chaos Approximations for L^2-Rndom Variables*. Sandia Report SAND2007-1262, 2007.

[14] Xiaoliang Wan and George Em Karniadakis, *An Adaptive Multi-element Generalized Polynomial Chaos Method for Stochastic Differential Equations*, J. Comput. Phys. 209 (2005) 617-642.

[15] Xiaoliang Wan and George Em Karniadakis, *Multi-element Generalized Polynomial Chaos for Arbitrary Probability Measures*, SIAM J. Sci. Comput. 28 (3) (2006) 901-928.

[16] Jean-Paul Berrut and Lloyd N. Trefethen, *Barycentric Lagrange Interpolation*, SIAM Rev. 46 (3) (2004) 501-517.

[17] R.V. Field Jr. and M. Grigoriu, *On the Accuracy of the Polynomial Chaos Approximation*, J. Comput. Phys. 209 (2005) 617-642.

[18] Nitin Agarwal and N.R. Aluru, *Weighted Smolyak Algorithm for Solution of Stochastic Differential Equations on Non-uniform Probability Measures*, Int. J. Numer. Meth. Engng. 85 (2011) 1365-1389.

[19] A. Klimke. *Uncertainty Modeling using Fuzzy Arithmetic and Sparse Grids*, PhD Thesis, Universität Stuttgart, Shaker Verlag, Aachen, 2006.

[20] A. Klimke, B. Wohlmuth, *Algorithm 847: spinterp: Piecewise Multilinear Hierarchical Sparse Grid Interpolation in MATLAB*, ACM Transactions on Mathematical Software 31 (2005).

[21] A. Klimke, *Sparse Grid Interpolation Toolbox – User's Guide*, IANS report 2006/001, University of Stuttgart, 2006.

[22] R. Sibson, *A brief description of natural neighbor interpolation (Chapter 2)*. In V. Barnett. Interpreting Multivariate Data. Chichester: John Wiley (1981) 21-36.

[23] R. Sibson, *A vector identity for the Dirichlet tessellaions.* Math. Proc. Cambridge Philos. Soc. 87 (1980) 151-155.

[24] Gerald Farin, *Surfaces over Dirichlet Tessellations*, Computer Aided Geometric Design 7 (1990) 281-292.

[25] B. Piper, *Properties of Local Coordinates Based on Dirichlet Tessellations*, Computing Suppl. 8 (1993) 227-239.

[26] Jean-Daniel Boissonnat and Julia Flötotto, *A coordinate system associated with points scattered on a surface*, Computer-Aided Design. 36 (2004) 161-174.

[27] Tom Bobach, Gerald Farin, Dianne Hansford, and Georg Umlauf, *Natural neighbor extrapolation using ghost points*, Computer-Aided Design. 41 (2009) 350-365.

[28] Tom M. Apostol, *Mathematical Analysis*, 2nd Edition. Addison-Wesley Publishing Co., Reading, Mass., 1974.

[29] Martin Bertram, Shirley E. Konkle, Hans Hagen, Bernd Hamann, and Kenneth I. Joy. Terrain Modeling Using Voronoi Hierarchies. In G. Farin, H. Hagen, and B. Hamann, editors, *Hierarchical and Geometrical Methods in Scientific Visualization.* Springer, 2003, pp. 89-97.