

STOCHASTIC COLLOCATION WITH KERNEL DENSITY ESTIMATION*

HOWARD C. ELMAN[†] AND CHRISTOPHER W. MILLER[‡]

Abstract. The stochastic collocation method has recently received much attention for solving partial differential equations posed with uncertainty, i.e., where coefficients in the differential operator, boundary terms or right-hand sides are random fields. Recent work has led to the formulation of an adaptive collocation method that is capable of accurately approximating functions with discontinuities and steep gradients. These methods, however, usually depend on an assumption that the random variables involved in expressing the uncertainty are independent with marginal probability distributions that are known explicitly. In this work we combine the adaptive collocation technique with kernel density estimation to approximate the statistics of the solution when the joint distribution of the random variables is unknown.

1. Problem Statement. Let (Ω, Σ, P) be a complete probability space with sample space Ω , σ -algebra $\Sigma \subset 2^\Omega$ and probability measure $P : \Sigma \rightarrow [0, 1]$. Let $D \subset \mathbb{R}^d$ be a d -dimensional bounded domain with boundary ∂D . We investigate partial differential equations (PDEs) of the form

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \omega; u) &= f(\mathbf{x}), & \forall \mathbf{x} \in D, \omega \in \Omega \\ \mathcal{B}(\mathbf{x}, \omega; u) &= g(\mathbf{x}), & \forall \mathbf{x} \in \partial D, \omega \in \Omega. \end{aligned} \tag{1.1}$$

Here \mathcal{L} is a partial differential operator with boundary operator \mathcal{B} , both of which can depend on the random parameter ω . As a consequence of the Doob-Dynkin lemma, it follows that u is also a random field, dependent on both the spatial location \mathbf{x} and the event ω . In order to work numerically with the expressions in (1.1), we must first represent the operators in terms of a finite number of random variables $\boldsymbol{\xi} = [\xi_1, \xi_2, \dots, \xi_M]^T$. This is often accomplished using a truncated Karhunen-Loève (KL) expansion [13]. If we denote $\Gamma = \text{Image}(\boldsymbol{\xi})$, then we can write (1.1) as

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \boldsymbol{\xi}; u) &= f(\mathbf{x}), & \forall \mathbf{x} \in D, \boldsymbol{\xi} \in \Gamma \\ \mathcal{B}(\mathbf{x}, \boldsymbol{\xi}; u) &= g(\mathbf{x}), & \forall \mathbf{x} \in \partial D, \boldsymbol{\xi} \in \Gamma. \end{aligned} \tag{1.2}$$

For a given realization of the random vector $\boldsymbol{\xi}$, the system (1.2) is a deterministic partial differential equation that can be solved using a deterministic solver. Throughout this paper we assume that D , \mathcal{L} , \mathcal{B} , f , and g are defined so that the above problem (1.2) is well posed for all values of $\boldsymbol{\xi} \in \Gamma$. In this paper we will explore several different sampling methods for solving the system (1.2).

One is typically interested in methods that allow statistical properties of u to be computed. If $\rho(\boldsymbol{\xi})$ denotes the joint probability density function of the random vector $\boldsymbol{\xi}$, then the k^{th} moment of the solution u is defined as

$$\mathbb{E}(u^k) = \int_{\Gamma} u^k \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}. \tag{1.3}$$

*This work was supported in part by the U.S. Department of Energy under grant DEFG0204ER25619 and the U.S. National Science Foundation under grants CCF0726017 and DMS1115317.

[†]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 (elman@cs.umd.edu).

[‡]Department of Applied Mathematics and Scientific Computation, University of Maryland, College Park, MD 20742 (cmiller@math.umd.edu).

One may also be interested in computing probability distributions associated with u , for example $P(u(\mathbf{x}, \boldsymbol{\xi}) \geq c)$.

Several methods have been developed for computing approximations to the random field u and the associated statistical quantities. The most widely known is the Monte-Carlo method, where the desired statistics are obtained by repeatedly sampling the distribution of $\boldsymbol{\xi}$, solving each of the resulting deterministic PDEs, and then estimating the desired quantities by averaging. Recently, much attention has been paid to alternative approaches such as the stochastic Galerkin and stochastic sparse grid collocation methods [1, 7, 8, 17, 18, 23]. These methods typically approximate the solution u as a high-degree multivariate polynomial in $\boldsymbol{\xi}$. If this approximation is denoted $u_p(\mathbf{x}, \boldsymbol{\xi})$, then the error $u - u_p$ can be measured in terms of an augmented Sobolev norm

$$\|\cdot\|_{L^2_P;V} = \left(\int_{\Omega} \|\cdot\|_V^2 dP \right)^{\frac{1}{2}}. \quad (1.4)$$

Here V is an appropriate Sobolev space that depends on the spatial component of the problem and $\|\cdot\|_V$ is the norm over this space. It can be shown that as the total degree of the polynomial approximation is increased, the error in the above norm, $\|u - u_p\|_{L^2_P;V}$, decays very rapidly provided that the solution u is sufficiently smooth in $\boldsymbol{\xi}$ [17]. If u is not sufficiently smooth then the convergence of these methods can stall or they may not converge at all [14]. Several methods have been proposed for treating problems that are discontinuous in the stochastic space. One approach partitions the stochastic space into elements and approximates the solution locally within elements by polynomials, continuous on the domain [2, 21]. Another approach is to use a hierarchical basis method developed in [12], which approximates u using a hierarchical basis of piecewise linear functions defined on a sparse grid. This idea was used with stochastic collocation in [14] where the sparse grid is refined adaptively using an a posteriori error estimator.

If the truncated Karhunen-Loève expansion is used to express \mathcal{L} and \mathcal{B} , then the random variables $\xi_1, \xi_2, \dots, \xi_M$ have zero mean and are uncorrelated [13]. It is frequently assumed that the random variables are independent and that their marginal density functions $\rho_i(\xi_i)$ are known explicitly. In this case the joint density function is simply the product of the marginal densities $\rho(\boldsymbol{\xi}) = \prod_{k=1}^M \rho_k(\xi_k)$. This assumption simplifies the evaluation of the moments of the solution since the multidimensional integral in (1.3) can be written as the product of one-dimensional integrals. It is not the case, however, that uncorrelated random variables are necessarily independent, and in the worst case the support of the product of the marginal densities may contain points that are not in the support of the true joint density. Thus, it may not be appropriate to define the joint density function as the product of the marginal density functions. See [9] for further discussion of this point. In this paper we explore a method for approximating the statistics of the solution u when an explicit form of the joint distribution is not available and we only have access to a finite number of samples of the random vector $\boldsymbol{\xi}$. In particular, we are able to treat the case where information on the parameters of the problem is only available in the form of experimental data. The method works by constructing an approximation $\hat{\rho}(\boldsymbol{\xi})$ to the joint probability distribution $\rho(\boldsymbol{\xi})$ using kernel density estimations [19]. This construction is then combined with an adaptive collocation strategy similar to the one derived in [14] to compute an approximation to the random field u . Moments can then be efficiently evaluated by integrating this approximation with respect to

the approximate probability measure $\hat{\rho}(\boldsymbol{\xi})$.

The remainder of this paper proceeds as follows. Section 2 discusses the adaptive collocation method in [14]. Section 3 presents an overview of the kernel density estimation technique used for approximating the unknown distribution of $\boldsymbol{\xi}$. Section 4 presents the method developed in this paper for approximating solutions to problems of the form (1.2). An error bound for the method is given in Section 4.1, and Section 4.2 presents techniques for extracting solution statistics. Section 5 presents the results of numerical experiments showing the performance of the new method and comparing this performance with that of the Monte Carlo method. Finally in Section 6 we draw some conclusions.

2. The Adaptive Collocation Method. Collocation methods work by solving the equation (1.2) for a finite number of pre-determined parameters $\{\boldsymbol{\xi}^{(1)}, \dots, \boldsymbol{\xi}^{(N_c)}\}$ using a suitable deterministic solver. The solutions at each sample point are then used to construct an interpolant to the solution for arbitrary choices of the random vector $\boldsymbol{\xi}$. We denote such an approximation generally as $\mathcal{A}(u)(\boldsymbol{\xi})$. Collocation methods were first used for solving PDEs with random coefficients in [1]. The interpolant was formed using a Lagrangian polynomial basis defined on tensor product grids. The cardinality of these grids is exponential in the dimension of the random vector so that this method is not viable for problems with high-dimensional random inputs. Sparse grid collocation methods were developed in [23] and an error analysis of the method was presented in [17]. These methods use the Smolyak interpolation formula [20] to construct a high-order polynomial interpolant using many fewer points than the full tensor grid. A refinement of this method for problems where the solution depends on the parameters in an anisotropic manner was presented in [18]. For all of these methods, the solution random field is expressed globally as a polynomial in the random vector $\boldsymbol{\xi}$. These methods are therefore only useful when the random field u is sufficiently regular in $\boldsymbol{\xi}$.

An adaptive collocation method was developed in [14]. This method is designed to compute approximations of random fields that possess discontinuities or strong gradients, and for which the image set Γ is bounded.¹ In the following, we present an overview of this method and our proposed modifications. To simplify the presentation we describe the case of a function u defined by a single random parameter whose image is a subset of $[0, 1]$. This can be generalized in a straightforward manner to a function defined by M parameters with image contained in any M -dimensional hypercube. Define

$$m_i = \begin{cases} 1 & \text{if } i = 1, \\ 2^{i-1} + 1 & \text{if } i > 1, \end{cases} \quad (2.1)$$

$$\xi_j^i = \begin{cases} \frac{j-1}{m_i-1} & \text{for } j = 1, \dots, m_i, \text{ if } m_i > 1, \\ 0.5 & \text{for } j = 1, \text{ if } m_i = 1. \end{cases} \quad (2.2)$$

For $i = 1, 2, \dots$, we have that $\theta^i = \{\xi_j^i\}_{j=1}^{m_i}$ consists of m_i distinct equally spaced points on $[0, 1]$. We also have that $\theta^i \subset \theta^{i+1}$. Since these points are equidistant, the use of global polynomial interpolation as in [23] is not appropriate due to the Runge phenomenon. We make no assumptions on the smoothness of u ; for example, it may contain singularities that global polynomial approximations will not resolve. To address these issues, a hierarchical basis of piecewise linear functions is used to

¹For unbounded Γ , interpolation is carried out on a bounded subset of Γ , see e.g. [22].

construct the interpolant. Define $\theta^0 = \emptyset$ and $\Delta\theta^i = \theta^i \setminus \theta^{i-1}$. Note that $|\Delta\theta^i| = m_i - m_{i-1}$. Let the members of $\Delta\theta^i$ be denoted $\{\xi_j^{\Delta^i}\}_{j=0}^{|\Delta\theta^i|-1}$. The hierarchical basis is defined on the interval $[0, 1]$ as

$$a_0^1(\xi) = 1 \quad (2.3)$$

$$a_j^i(\xi) = \begin{cases} 1 - (m_i - 1)|\xi - \xi_j^{\Delta^i}| & \text{if } |\xi - \xi_j^{\Delta^i}| < 1/(m_i - 1), \\ 0 & \text{otherwise,} \end{cases} \quad (2.4)$$

for $i > 1$ and $j = 0, \dots, |\Delta\theta^i| - 1$. These functions are piecewise linear and have the




FIG. 2.1. The hierarchical basis functions for $i = 1, 2, 3$.

property that $a_j^i(\xi_k^{\Delta^i}) = \delta_{jk}$, and $a_j^i(\xi_k^s) = 0$ for all $s < i$. Note that there is a binary tree structure on the nodes in θ^i . That is, we can define the set of children of a point $\xi_j^{\Delta^i}$ as

$$child(\xi_j^{\Delta^i}) = \begin{cases} \{\xi_j^{\Delta^{i+1}}\} & \text{if } i = 2 \\ \{\xi_{2j}^{\Delta^{i+1}}, \xi_{2j+1}^{\Delta^{i+1}}\} & \text{otherwise.} \end{cases} \quad (2.5)$$

We also denote the parent of a point in this tree as $par(\xi_j^{\Delta^i})$.

Algorithm 1 defines an interpolation scheme using the hierarchical basis functions. The quantities $\{w_j^k\}$ are referred to as the hierarchical surplus. They represent the correction to the interpolant $\mathcal{A}_{i-1}(u)$ at the points in $\Delta\theta_i$. For functions with values that vary dramatically at neighboring points, the hierarchical surpluses $\{w_j^k\}$ remain large for several iterations. This provides us with a natural error indicator as well as a convergence criterion for the method, whereby we require that the largest hierarchical surplus be smaller than a given tolerance. The hierarchical surpluses also provide a

Algorithm 1 Interpolation With Hierarchical Basis Functions

Define $\mathcal{A}_0(u)(\xi) = 0$.
 Define $k = 1$
repeat
 Construct $\Delta\theta^k$
 Evaluate $u(\xi_j^{\Delta k}) \forall \xi_j^{\Delta k} \in \Delta\theta^k$
 $w_j^k = u(\xi_j^{\Delta k}) - \mathcal{A}_{k-1}(u)(\xi_j^{\Delta k}) \forall \xi_j^{\Delta k} \in \Delta\theta^k$
 Define $\mathcal{A}_k(u)(\xi) = \sum_{i=1}^k \sum_{j=0}^{|\Delta\theta^i|-1} w_j^i a_j^i(\xi)$.
 $k = k + 1$
until $\max(|w_j^{k-1}|) < \tau$

mechanism to implement adaptive grid refinement. The grid is adaptively refined at points with large hierarchical surpluses. For such a point, its children are added to the next level of the grid. Algorithm 2 defines such an adaptive interpolation algorithm that is similar to the one appearing in [14]. The interpolation error associated with

Algorithm 2 Adaptive Interpolation With Hierarchical Basis Functions

Define $\mathcal{A}_0(u)(\xi) = 0$.
 Define $k = 1$
 Initialize $\Delta\theta_{adaptive}^1 = \theta^1$.
repeat
 $\Delta\theta_{adaptive}^{k+1} = \emptyset$
 for $\xi_j^{\Delta k} \in \Delta\theta_{adaptive}^k$ **do**
 Evaluate $u(\xi_j^{\Delta k})$
 $w_j^k = u(\xi_j^{\Delta k}) - \mathcal{A}_{k-1}(u)(\xi_j^{\Delta k})$
 if $\|w_j^k\| > \tau$ **then**
 $\Delta\theta_{adaptive}^{k+1} = \Delta\theta_{adaptive}^{k+1} \cup \text{child}(\xi_j^{\Delta k})$
 end if
 end for
 Define $\mathcal{A}_k(u)(\xi) = \sum_{i=1}^k \sum_j w_j^i a_j^i(\xi)$.
 $k = k + 1$
until $\max(\|w_j^{k-1}\|) < \tau$

this method is shown by numerical experiments in [14] to be significantly smaller than the bound $\mathcal{O}(|\theta^k|^{-2} \log(|\theta^k|^{3(M+1)}))$ presented in [11] for both smooth functions and examples that contain steep gradients or discontinuities.

This method can be generalized in a straightforward way to functions defined on $[0, 1]^M$. All that is needed is to define a multidimensional hierarchical basis set and a method for generating the children of a given grid point. The multidimensional hierarchical basis consists of tensor products of the one-dimensional hierarchical basis functions. Given $\mathbf{i} = [i_1, \dots, i_M] \in \mathbb{N}^M$ and $\mathbf{j} = [j_1, \dots, j_M] \in \mathbb{N}^M$, let

$$a_{\mathbf{j}}^{\mathbf{i}}(\boldsymbol{\xi}) = a_{j_1}^{i_1}(\xi_1) \otimes \dots \otimes a_{j_M}^{i_M}(\xi_M). \quad (2.6)$$

We can define the multidimensional interpolation grids by

$$\begin{aligned} \theta_1 &= [0.5, 0.5, \dots, 0.5] \\ \text{child}(\boldsymbol{\xi}_j^{\Delta i}) &= \{\boldsymbol{\xi} | \exists! j \in 1, \dots, M \text{ s.t. } [\xi_1, \dots, \xi_{j-1}, \text{par}(\xi_j), \xi_{j+1}, \dots, \xi_M] = \boldsymbol{\xi}_j^{\Delta i}\}. \end{aligned} \quad (2.7)$$

From this we can see that each grid point has at most $2M$ children.

This method can be used to approximate the solutions to (1.2) by applying a suitable deterministic solver to the equations at collocation points $\boldsymbol{\xi}_j^{\Delta i}$. We can then construct an interpolant of u , $\mathcal{A}_k(u)$ using the formula in Algorithm 2. In principle, the expected value of u can be approximated by

$$\mathbb{E}(u) \approx \int_{\Gamma} \mathcal{A}_k(u) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi} = \sum_i \sum_j w_j^i \int_{\Gamma} a_j^i(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}, \quad (2.8)$$

although in the cases under discussion ρ will not be known explicitly. Even in the case where ρ is known explicitly and can be expressed as the product of univariate functions, the integral in (2.8) can still be difficult to calculate when it is of high dimension.

3. Kernel Density Estimation. Let $K(\boldsymbol{\xi})$ be a function satisfying the following conditions:

$$\begin{aligned} \int_{\mathbb{R}^M} K(\boldsymbol{\xi}) d\boldsymbol{\xi} &= 1, \\ \int_{\mathbb{R}^M} K(\boldsymbol{\xi}) \boldsymbol{\xi} d\boldsymbol{\xi} &= 0, \\ \int_{\mathbb{R}^M} K(\boldsymbol{\xi}) \|\boldsymbol{\xi}\|^2 d\boldsymbol{\xi} &= k_2 < \infty, \\ K(\boldsymbol{\xi}) &\geq 0, \end{aligned} \quad (3.1)$$

where $\|\boldsymbol{\xi}\|$ is the Euclidean norm of the M -dimensional vector $\boldsymbol{\xi}$. Let $\boldsymbol{\xi}^{(1)}, \boldsymbol{\xi}^{(2)}, \dots, \boldsymbol{\xi}^{(N)}$ be N independent realizations of the random vector $\boldsymbol{\xi}$. The kernel density approximation to the joint distribution of $\boldsymbol{\xi}$ is given by

$$\hat{\rho}(\boldsymbol{\xi}) = \frac{1}{Nh^M} \sum_{k=1}^N K\left(\frac{\boldsymbol{\xi} - \boldsymbol{\xi}^{(i)}}{h}\right), \quad (3.2)$$

where h is a user-defined parameter called the bandwidth. It is straightforward to verify that the function $\hat{\rho}$ defined above satisfies the conditions for being a probability density function. The main challenge here lies in the selection of an appropriate value for h . If h is chosen to be too large then the resulting estimate is said to be over-smoothed and important features of the data may be obscured. If h is chosen to be too small then the resulting estimate is said to be under-smoothed and the approximation may contain many spurious features not present in the true distribution. Figure 3.1 shows kernel density estimates of a bimodal distribution for a small and large value of h . The over-smoothed estimate does not detect the bimodality of the data whereas the under-smoothed estimate introduces spurious oscillations into the estimate.

One method for specifying h is to choose the value that minimizes the approximate mean integrated square error (*AMISE*). For a given value of h , the *AMISE*

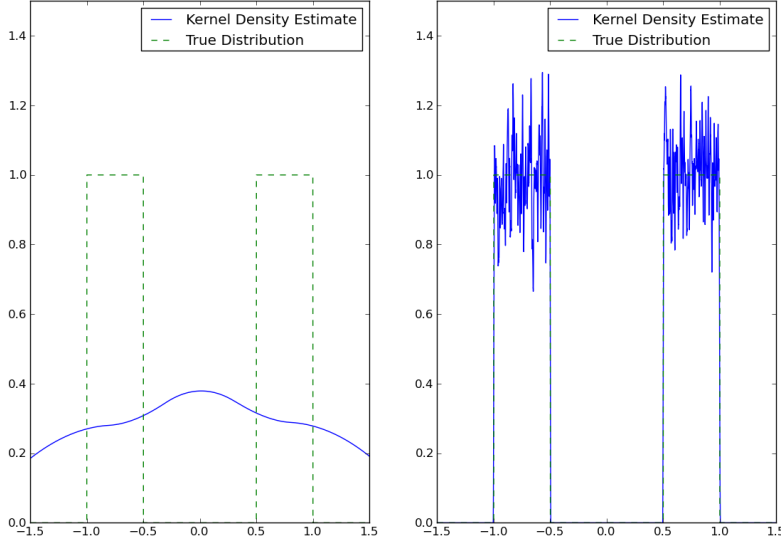


FIG. 3.1. Under-smoothed kernel density estimate (left) and over-smoothed (right).

is given by

$$AMISE(h, N) = \frac{1}{4}h^4\alpha^2 \int_{\mathbb{R}^M} (\Delta\rho(\boldsymbol{\xi}))^2 d\boldsymbol{\xi} + N^{-1}h^{-M}\beta, \quad (3.3)$$

where

$$\alpha = \int_{\mathbb{R}^M} \|\boldsymbol{\xi}\|_1^2 K(\boldsymbol{\xi}) d\boldsymbol{\xi}, \quad \beta = \int_{\mathbb{R}^M} K(\boldsymbol{\xi})^2 d\boldsymbol{\xi}, \quad (3.4)$$

and Δ here denotes the Laplace operator [19]. From this expression the optimal value of h can be derived as [19]

$$h_{opt}^{M+4} = M\beta\alpha^{-2} \left\{ \int (\Delta\rho(\boldsymbol{\xi}))^2 d\boldsymbol{\xi} \right\}^{-1} N^{-1}. \quad (3.5)$$

It can be shown that the optimal bandwidth is of magnitude $\mathcal{O}(N^{-1/(M+4)})$ as the number of samples N increases. If the optimal value of h is used it can also be shown that the $AMISE$ decays like $\mathcal{O}(N^{-\frac{4}{4+M}})$.

For numerical computations, choosing h to minimize the $AMISE$ is impractical since it requires a priori knowledge of the exact distribution. Many techniques have been proposed for choosing the smoothing parameter h without a priori knowledge of the underlying distribution, including least-squares cross-validation and maximum likelihood cross-validation [19]. In the numerical experiments below we employ maximum likelihood cross-validation (MLCV). This method proceeds as follows. Given a finite set of samples, $\boldsymbol{\xi}^{(1)}, \boldsymbol{\xi}^{(2)}, \dots, \boldsymbol{\xi}^{(N)}$, of the random vector $\boldsymbol{\xi}$, define

$$\hat{\rho}_{-i}(\boldsymbol{\xi}) = \frac{1}{Nh^M} \sum_{k=1, k \neq i}^N K\left(\frac{\boldsymbol{\xi} - \boldsymbol{\xi}^{(k)}}{h}\right) \quad (3.6)$$

to be the kernel density estimate constructed by omitting the i^{th} sample. The maximum likelihood cross-validation method is to choose h that maximizes

$$CV(h) \equiv \frac{1}{N} \sum_{i=1}^N \log(\hat{\rho}_{-i}(\boldsymbol{\xi}^{(i)})). \quad (3.7)$$

Note that this value of h only depends on the data. The intuition behind this method is that if we are given an approximation to the true density based on $N - 1$ samples and we draw another sample, then the approximate density should be large at this new sample point. In the numerical experiments described below, we solved this optimization problem using Brent's method [4]. The asymptotic cost of evaluating (3.7) is $\mathcal{O}(N^2)$. Thus as the number of samples grows large this method can become costly. In this case one typically only uses a randomly selected subset of the samples to evaluate (3.7) [10]. In the numerical experiments described below, we observed that for the sample sizes used, the cost of this optimization was significantly lower than the cost of repeatedly solving the algebraic systems of equations that arise from the spatial discretization of the PDE (1.2).

In [19] it is shown that the choice of kernel does not have a strong effect on the error associated with kernel density estimation. In our experiments we use the multivariate Epanechnikov kernel

$$K(\boldsymbol{\xi}) = \left(\frac{3}{4}\right)^M \prod_{i=1}^M (1 - \xi_i^2) \mathbf{1}_{\{-1 \leq \xi_i \leq 1\}}. \quad (3.8)$$

This kernel is frequently used in the case of univariate data as it minimizes the asymptotic mean integrated square error over all choices of kernels satisfying (3.1). It also has the advantage that it is compactly supported. This causes the approximate density function $\hat{\rho}$ to be compactly supported, which is important in assuring the well-posedness of some stochastic partial differential equations.

4. Adaptive Collocation With KDE Driven Grid Refinement. The interpolation method in [14] distributes interpolation nodes so that discontinuities and steep gradients in the solution function are resolved; however the method does not take into account how significant a given interpolation node is to the statistics of the solution function since the refinement process does not depend on ρ . The kernel density estimate described above can also be used to drive refinement of the adaptive sparse grid in Algorithm 2. The algorithm we propose is as follows. First construct an estimate $\hat{\rho}$ to the true density ρ using a finite number of samples $\{\boldsymbol{\xi}^{(i)}\}_{i=1}^N$. Second, replace the refinement criterion in Algorithm 2 with

$$|w_j^k| \hat{\rho}(\boldsymbol{\xi}_j^{\Delta k}) > \tau. \quad (4.1)$$

A similar approach is used in [15] to drive the refinement. However in that study it is again assumed that one has access to an explicit form of the joint density function. With the refinement criterion (4.1), the grid is only adaptively refined at points near the data $\{\boldsymbol{\xi}^{(i)}\}_{i=1}^N$ since the kernel density estimate is only supported near the samples. In the sequel we refer to this proposed method, i.e., Algorithm 2 with refinement criterion (4.1), as *adaptive KDE collocation*. The remainder of this section is divided into two parts. In Section 4.1 we present interpolation error estimates associated with adaptive KDE collocation and in Section 4.2 we present methods for approximating the solution statistics of the random field u . Note that throughout this discussion we can ignore the spatial component of the problem.

4.1. Error analysis of adaptive KDE collocation. For simplicity we present the results for the case where the problem only depends on a single parameter and that interpolation is carried out on $[0, 1]$. Extension of the argument to multi-parameter problems defined on an arbitrary hypercube is straightforward. Also we ignore the spatial component of the problem as it has no effect on the discussion of the errors resulting from the discretization of the stochastic portion of the problem. Assume that $\mathcal{A}_k(u)$ is an interpolant generated using adaptive KDE collocation with tolerance τ . Let $\hat{\rho}$ be the kernel density estimate used in computing \mathcal{A}_k and let $\hat{\Gamma}$ be the support of $\hat{\rho}$. Let $\mathcal{A}_k^{complete}(u)$ be the interpolant constructed by Algorithm 1 with grid points $\Delta\theta^k = \{\xi_j^{\Delta^i}\}$ and set of hierarchical surpluses $\{w_j^i\}$ at those grid points. By definition, $\Delta\theta_{adaptive}^k \subset \Delta\theta^k$. Define $\Delta\theta_{remaining}^k = \Delta\theta^k \setminus \Delta\theta_{adaptive}^k$. Then if $\xi_j^{\Delta^i} \in \Delta\theta_{remaining}^k$, it follows from (4.1) that $|w_j^i \hat{\rho}(\xi_j^{\Delta^i})| \leq \tau$. We can bound the difference between u and $\mathcal{A}_k(u)$ on $\hat{\Gamma}$ as

$$\begin{aligned} \|(u - \mathcal{A}_k(u))\hat{\rho}\|_{L_\infty(\hat{\Gamma})} &\leq \left\| \frac{\hat{\rho}}{\hat{\rho}} \right\|_{L_\infty(\hat{\Gamma})} \underbrace{\|(u - \mathcal{A}_k^{complete}(u))\hat{\rho}\|_{L_\infty(\hat{\Gamma})}}_{\epsilon_1} + \\ &\quad \underbrace{\|(\mathcal{A}_k^{complete}(u) - \mathcal{A}_k(u))\hat{\rho}\|_{L_\infty(\hat{\Gamma})}}_{\epsilon_2}. \end{aligned} \quad (4.2)$$

The term ϵ_1 is the interpolation error associated with piecewise multilinear approximation on a full grid. This case is studied in [11]. The interpolation error is bounded by

$$\|u - \mathcal{A}_k^{complete}(u)\|_{L_\infty(\Gamma)} = \mathcal{O}(|\Delta\theta^k|^{-2} |\log_2(|\Delta\theta^k|)|^{3(M-1)}) \quad (4.3)$$

Since $\hat{\rho}$ is bounded it follows that the bound on ϵ_1 decays at the same rate.

Bounding ϵ_2 depends on counting the points in $\Delta\theta_{remaining}^k$ and using the fact that at those points $|w_j^i \hat{\rho}| \leq \tau$. We have that

$$\|(\mathcal{A}_k^{complete}(u) - \mathcal{A}_k(u))\hat{\rho}\|_{L_\infty(\Gamma)} \leq \sum_{\Delta\theta_{remaining}^k} |w_j^i| \|a_j^i(\xi)\hat{\rho}(\xi)\|_{L_\infty(\Gamma)}. \quad (4.4)$$

Expanding $\hat{\rho}$ in a Taylor series around $\xi_j^{\Delta^i}$ and noting that $a_j^i(\xi)\hat{\rho}(\xi)$ is only supported on an interval of size $\frac{1}{2^i}$ gives

$$\begin{aligned} \|(\mathcal{A}_k^{complete}(u) - \mathcal{A}_k(u))\hat{\rho}\|_{L_\infty(\Gamma)} &\leq \sum_{\Delta\theta_{remaining}^k} |w_j^i \hat{\rho}(\xi_j^{\Delta^i})| + |w_j^i| \frac{\|\hat{\rho}'\|_{L_\infty(\Gamma)}}{2^i} \\ &\leq \tau |\Delta\theta_{remaining}^k| + \sum_{\Delta\theta_{remaining}^k} |w_j^i| \frac{\|\hat{\rho}'\|_{L_\infty(\Gamma)}}{2^i}. \end{aligned} \quad (4.5)$$

The sums here are over all \mathbf{i}, \mathbf{j} such that $\xi_j^{\Delta^i} \in \Delta\theta_{remaining}^k$. For decreasing τ , the number of points in $\Delta\theta_{remaining}^k$ decreases, since more points are locally refined and those points that remain in $\Delta\theta_{remaining}^k$ for large k correspond to basis functions with very small support. If τ is chosen to be small and k is allowed to grow so that the refinement criterion (4.1) is satisfied at every leaf node, the term ϵ_2 will converge to zero.

4.2. Estimation of Solution Statistics. Computation of the moments of the solution via the methods presented in [1, 2, 8, 14, 17, 23] all require that the joint density function ρ be explicitly available in order to evaluate the integral $\int_{\Gamma} \hat{u}(\mathbf{x}, \boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}$ where \hat{u} is an approximation to u computed by either the stochastic Galerkin method [2, 8] or by the stochastic collocation method [1, 14, 17, 23]. In practice this may be an unrealistic assumption since we often only have access to a finite sample from the distribution of $\boldsymbol{\xi}$. This section describes two ways of approximating the solution statistics when only a random sample from the distribution of $\boldsymbol{\xi}$ is available. The first is the well-known Monte-Carlo method [16]; the second is a variant of the Monte-Carlo predictor method presented in [22].

Given a random field $u(x, \boldsymbol{\xi})$ and a finite number of samples $\{\boldsymbol{\xi}^{(i)}\}_{i=1}^N$, the Monte-Carlo method approximates the mean of u by the sample mean

$$\mathbb{E}(u)(\mathbf{x}) \approx \frac{1}{N} \sum_{i=1}^N u(\mathbf{x}, \boldsymbol{\xi}^{(i)}) \equiv \bar{u}(\mathbf{x}). \quad (4.6)$$

This method has the advantage that the convergence is independent of the dimension of the random parameter. The error in the expected value can be approximated by first noting that the estimate is unbiased,

$$Bias_{MC} = \mathbb{E}(u)(\mathbf{x}) - \mathbb{E} \left(\frac{1}{N} \sum_{i=1}^N u(\mathbf{x}, \boldsymbol{\xi}^{(i)}) \right) = 0, \quad (4.7)$$

and that

$$Var(\bar{u}(\mathbf{x})) = \frac{Var(u(\mathbf{x}, \boldsymbol{\xi}))}{N}, \quad (4.8)$$

where $Var(\bar{u}(\mathbf{x}))$ is the variance of the sample mean. An application of Chebyshev's inequality then gives a standard probabilistic estimate, that for $a > 0$,

$$P \left(\left| \mathbb{E}(u)(\mathbf{x}) - \frac{1}{N} \sum_{i=1}^N u(\mathbf{x}, \boldsymbol{\xi}^{(i)}) \right| \geq a \right) \leq \frac{Var(u)}{Na^2}. \quad (4.9)$$

Note that a factor of 2 error reduction requires an increase of the sample size by a factor of 4. This slow rate of convergence is often cited as the chief difficulty in using the Monte-Carlo method [1, 8]. It is also important to note that this bound is probabilistic in nature and that it is possible for the Monte-Carlo method to perform much worse (or much better) than expected. For a fixed choice of the quantity on the left hand side of (4.9), which we call P here, say $P = .05$, we have that

$$a \leq \sqrt{\frac{Var(u)}{.05N}}, \quad (4.10)$$

and from this we can conclude with 95% percent confidence that the Monte-Carlo estimate is bounded by $\sqrt{\frac{Var(u)}{.05N}}$. Smaller values of P lead to looser bounds but greater confidence in those bounds.

The method presented in [22] is to construct an approximation \hat{u} of the solution function in the stochastic space using conventional sparse grid collocation and then, given a finite number of samples $\{\boldsymbol{\xi}^{(i)}\}_{i=1}^N$, to approximate the expected value by

$$\mathbb{E}(u)(x) \approx \frac{1}{N} \sum_{i=1}^N \hat{u}(\mathbf{x}, \boldsymbol{\xi}^{(i)}). \quad (4.11)$$

Instead of using conventional sparse grid collocation, we construct an approximation \hat{u} using the adaptive KDE collocation method. Assuming that one has already constructed the interpolant, computation of the expected value can be carried out very quickly this way since the interpolant is simple to evaluate. Note also that while the standard Monte-Carlo method was used to evaluate (4.11), adaptive KDE collocation is also compatible with other sampling methods such as quasi-Monte Carlo [5] and multilevel Monte-Carlo [3, 6]. In the case of quasi-Monte Carlo, the sample points used in (4.11) are simply chosen to be the quasi-Monte Carlo sample points, and in the case of multilevel Monte-Carlo an expression similar to (4.11) is computed at each level of the computation. We expect that combining adaptive KDE collocation with either of these alternative sampling strategies would yield combined benefits; we do not explore this issue here.

The error associated with this method separates into two terms as follows,

$$\begin{aligned} |\epsilon_{sparse}| &= |\mathbb{E}(u)(\mathbf{x}) - \frac{1}{N} \sum_{i=1}^N \mathcal{A}(u)(\mathbf{x}, \boldsymbol{\xi}^{(i)})| \\ &\leq |\mathbb{E}(u)(\mathbf{x}) - \frac{1}{N} \sum_{i=1}^N u(\mathbf{x}, \boldsymbol{\xi}^{(i)})| + |\frac{1}{N} \sum_{i=1}^N (u(\mathbf{x}, \boldsymbol{\xi}^{(i)}) - \mathcal{A}(u)(\mathbf{x}, \boldsymbol{\xi}^{(i)}))| \\ &= \epsilon_{MC} + \epsilon_{interp}. \end{aligned} \tag{4.12}$$

The first term is statistical error and depends only on the number of samples taken and the variance of u , and decays according to (4.9). The second term is the interpolation error and is bounded since the infinity norm of the interpolation error is bounded in the neighborhood of the sample points using (4.2).

Given N samples of $\boldsymbol{\xi}$, evaluation of (4.6) requires N evaluations of the random field u . In the case where u is defined by a system such as (1.2), this requires N solutions of a discrete PDE. In contrast, evaluation of (4.11) requires N_{interp} evaluations of u to construct $\mathcal{A}(u)$ and then it requires N evaluations of $\mathcal{A}(u)$. The relative computational efficiency of (4.11) then depends on two factors: first, whether an accurate interpolant $\mathcal{A}(u)$ can be constructed using $N_{interp} \ll N$ function evaluations, and second, whether the cost of evaluating $\mathcal{A}(u)$ is significantly less than the cost of evaluating u . The first condition, as shown by (4.3), depends on the dimension of the problem as well as the number of samples we have access to. For most problems of interest the second condition is satisfied in that it is much less expensive to evaluate a piecewise polynomial than it is to solve a discrete algebraic system associated with a complex physical model. Note that in order for ϵ_{interp} to be small the interpolation error only needs to be small near the sample points. For adaptive KDE collocation the kernel density estimate is designed to make the interpolant more accurate in the neighborhoods of these points by indicating where large clusters of points are located.

5. Numerical Experiments. In this section we assess the performance of adaptive KDE collocation applied to several test problems. We aim to measure quantitatively the two terms in the estimate (4.2) and to compare the computational efficiency of our method with the Monte-Carlo method.

5.1. Example 1: Interpolation of a Highly Oscillatory Function. Before exploring our main concern, the solution of PDEs with stochastic coefficients, we first examine the utility of adaptive collocation for performing a simpler task, to interpolate a scalar-valued function whose argument is a random vector. We use adaptive KDE

collocation to construct an approximation to the function

$$u(\boldsymbol{\xi}) = \begin{cases} \prod_{k=1}^M |\xi_k| \sin(1/\xi_k) & \text{if } \xi_k \neq 0 \\ 0 & \text{otherwise,} \end{cases} \quad (5.1)$$

where $\boldsymbol{\xi}$ is a random variable uniformly distributed over the set $[-1, -0.5]^M \cup [0.5, 1]^M$. Figure 5.1 shows a plot of the function $u(\xi)$ for the single parameter case. The density of $\boldsymbol{\xi}$ is given explicitly by

$$\rho(\xi) = 2^{M-1} \mathbf{1}_{[-1, -0.5]^M \cup [0.5, 1]^M}. \quad (5.2)$$

The function u is everywhere continuous but infinitely oscillatory along each axis of $\boldsymbol{\xi}$. The axes however are not contained in the support of ρ so the oscillations do not have any effect on the statistics of u with respect to the measure on $\boldsymbol{\xi}$. Algorithm 2 with the refinement criterion used in [14] would place many collocation points near the origin in an attempt to resolve the oscillatory behavior. Provided that the approximate density $\hat{\rho}$ is a good approximation to the true density, adaptive KDE collocation will only place collocation points near the support of ρ .




FIG. 5.1. $u(\xi) = |\xi| \sin(1/\xi)$.

In our experiments, the density estimate for each choice of M will be constructed from 5,000 samples of $\boldsymbol{\xi}$ with the bandwidth h chosen by maximum likelihood cross validation. For a given value of ξ let $|(u(\xi) - \mathcal{A}_k(u)(\xi))\rho(\xi)|$ be the interpolation error scaled by ρ . First we measure the scaled interpolation error at 500 equally spaced points on $[-1.5, 1.5]$ and use the maximum observed error as an estimate for the infinity norm of the error $\|(u(\xi) - \mathcal{A}_k(u)(\xi))\rho(\xi)\|_{L^\infty(\Gamma)}$ for the one-parameter (i.e. $M = 1$ in (5.1)) problem. We denote this estimate by $\|(u(\xi) - \mathcal{A}_k(u)(\xi))\rho(\xi)\|_{l^\infty}$. Figure 5.2 shows the interpolation error in the mesh-norm $\|\cdot\|_{\rho(\xi)}$. This norm only indicates the error on the support of ρ . Figure 5.2 shows that the interpolation error decays rapidly where the random variable ξ is supported. Figure 5.2 shows that adaptive KDE collocation converges significantly faster than Algorithm 2. The




FIG. 5.2. $\|(u(\xi) - \mathcal{A}_k(u)(\xi))\rho(\xi)\|_\infty$ versus the number of collocation points

reason is that Algorithm 2 places many points near the origin, attempting to resolve the oscillations. After a few initial global refinements of the grid the new method concentrates all of the new collocation points inside the support of ξ .² Figure 5.3 shows the collocation nodes used by the adaptive method with KDE driven refinement.

Now we examine the performance for the same task when u depends on multiple parameters in (5.1). Figure 5.4 shows the number of collocation points required as a function of the convergence criterion τ and the number of parameters. The figure shows that as the number of parameters is increased, the efficiency of the proposed method slows. This is due to the factor $\log_2(|\Delta\theta^k|)^{3(M-1)}$ appearing in the estimate (4.3). Note however that for any fixed value of M , the asymptotic interpolation error bound (4.3) decays faster than the Monte-Carlo error bound (4.9). The results in Section 5.3 indicate that the asymptotic bound (4.3) may be pessimistic for problems of interest.



5.2. Example 2: Two-parameter stochastic diffusion equation. Next, we use the method derived in section 4 to compute statistics associated with the solution to the stochastic diffusion equation

$$-\nabla \cdot (a(\mathbf{x}, \xi_1, \xi_2) \nabla u(\mathbf{x}, \xi_1, \xi_2)) = 1, \quad \forall \mathbf{x} \in D \quad (5.3)$$

$$u(\mathbf{x}, \xi_1, \xi_2) = 0, \quad \forall \mathbf{x} \in \partial D \quad (5.4)$$

where $D = [0, 1]^2$. The diffusion coefficient a is defined for this example as follows. Define the set $LL = \{\mathbf{x} : 0 < x_1, x_2 \leq 0.5\}$ and the set $UR = \{\mathbf{x} : 0.5 < x_1, x_2 < 1.0\}$. Let $\mathbb{1}_{LL}(\mathbf{x})$ and $\mathbb{1}_{UR}(\mathbf{x})$ be the indicator functions on LL and UR respectively. The

²Algorithm 2 with the refinement criterion (4.1) indicates that a node is not refined if $\hat{\rho} \|w_j^k\|$ is small. In practice however it is necessary to perform some initial global grid refinements to achieve a minimum level of resolution.

FIG. 5.3. $u(\xi)$ and the collocation points used in constructing approximate solutionFIG. 5.4. The tolerance τ vs the number of collocation points

diffusion coefficient is piecewise constant and is given by

$$a(\mathbf{x}, \xi_1, \xi_2) = 1 + \mathbb{1}_{LL}(\mathbf{x})\xi_1 + \mathbb{1}_{UR}(\mathbf{x})\xi_2. \quad (5.5)$$

Here ξ_1 and ξ_2 are assumed to be independently distributed log-normal random variables. The PDF of ξ_i for $i = 1, 2$ is given by

$$\rho_i(\xi_i) = \frac{1}{\xi_i \sqrt{2\pi\sigma^2}} e^{-\frac{(\log(\xi_i) - \mu)^2}{2\sigma^2}}, \quad (5.6)$$

with $\sigma = 1$ and $\mu = 2$. Since ξ_1 and ξ_2 are assumed to be independent, their joint distribution is given by

$$\rho(\xi_1, \xi_2) = \frac{1}{2\pi\xi_1\xi_2} e^{-\frac{(\log(\xi_1)-2)^2 - (\log(\xi_2)-2)^2}{2}}. \quad (5.7)$$

Note that ξ_1 and ξ_2 take on values in the range $(0, \infty)$. This, combined with the definition of the diffusion coefficient in (5.5) ensures that the diffusion coefficient will be positive at all points in D for all possible values of the random variables ξ_1 and ξ_2 . This is sufficient to ensure the well-posedness of (5.3) [1]. In the numerical experiments, interpolation was carried out on the domain $[1 \times 10^{-6}, 6]^2$. This computational domain contained all of the samples of (ξ_1, ξ_2) generated by the log-normal random number generator.

The method described above generates a set of collocation points in the stochastic space. At each of these points (5.3) must be solved by using a suitable deterministic solver. In this example the spatial discretization is accomplished using finite differences on a uniform 32×32 mesh. The discrete difference operators are formed using the five point stencil

$$\begin{bmatrix} & a(x, y + \frac{h_D}{2}, \xi_1, \xi_2) & \\ a(x - \frac{h_D}{2}, y, \xi_1, \xi_2) & a(x, y, \xi_1, \xi_2) & a(x + \frac{h_D}{2}, y, \xi_1, \xi_2) \\ & a(x, y - \frac{h_D}{2}, \xi_1, \xi_2) & \end{bmatrix}, \quad (5.8)$$

for $\mathbf{x} = [x, y]^T \in D$, and where h_D is the spatial discretization parameter. For this example the resulting linear systems are solved using a direct solver, although an iterative solver may also be used as in [7]. Although the spatial discretization of the problem introduces an additional source of error, it is known that the error resulting from the spatial discretization of the problem separates from the error associated with discretization of the stochastic component [1, 2]. Thus we can focus solely on the error introduced by interpolating in the stochastic space and by approximating the true joint density by a kernel density estimate.

First we proceed as in Section 5.1 and evaluate the interpolation error. Since the exact solution is not known we compute $\mathcal{A}(u)$ with a very tight error tolerance $\tau = 10^{-9}$. We treat this as an accurate solution and observe the decay in error for interpolants obtained using a looser error tolerance. For each interpolant, the kernel density estimate is derived from 5,000 samples of $\boldsymbol{\xi} = [\xi_1, \xi_2]$ where ξ_1 and ξ_2 are independently distributed log-normal random variables as described above. The bandwidth for the kernel density estimates is chosen using the maximum likelihood cross-validation method described in section 3.

Figure 5.5 shows the collocation points used for several values of the error tolerance τ . Comparing these with the contour plot of the true joint density function in Figure 5.6, it can be seen that the method is concentrating collocation points in regions where the estimated joint PDF is large. Thus the method is only devoting resources towards computing an accurate interpolant in regions that are significant to the statistics of u . Figure 5.7 shows the interpolation error as a function of the number of collocation points. Since an exact solution to (5.3) is not available we treat the solution obtained by using the method with $\tau = 10^{-10}$ as an exact solution. As opposed to the first example, the solution u here depends on both the spatial location and the value of the random parameter. We report the error in the discrete norm $\|\cdot\|_{l^2(D) \times l^\infty(\Gamma)}$, where the space $l^2(D)$ consists of square summable mesh-functions defined on the spatial grid and $l^\infty(\Gamma)$ consists of bounded mesh-functions defined on

a 500×500 uniform grid on Γ . Figure 5.7 shows that the interpolation error decays quickly for the two parameter problem. The apparent slowdown in convergence rate is attributable to the fact that the exact solution is not available and the error is being measured with respect to an approximate solution.




FIG. 5.5. Collocation points for various values of the error tolerance τ

5.3. High-dimensional stochastic diffusion. We now examine the performance of adaptive KDE collocation for evaluating the statistics of a random field that depends on a large number of parameters. The problem is given by

$$-\frac{d}{dx}(a_M(x, \boldsymbol{\xi}) \frac{d}{dx} u(x, \boldsymbol{\xi})) = 1, \quad \forall x \in (0, 1) \quad (5.9)$$

$$u(0, \boldsymbol{\xi}) = u(1, \boldsymbol{\xi}) = 0. \quad (5.10)$$

The diffusion coefficient a_M is defined for even M by

$$a_M = \mu + \sum_{k=0}^{M/2-1} \lambda_k (\xi_{2k} \cos(2\pi kx) + \xi_{2k+1} \sin(2\pi kx)), \quad (5.11)$$

where $\lambda_k = \exp(-k)$, $\mu = 3$ and ξ_k is uniformly distributed on $[0, 1]$. The problem (5.9) is well posed on the image of $\boldsymbol{\xi}$. Experimental results for these problems are shown in Tables 5.1 (for $M = 4$ random variables), 5.2 ($M = 10$), and 5.3 ($M = 20$). The contents of the tables are as follows.

First, for each M , we performed a Monte-Carlo simulation with several choices of number of samples N . This sample size is shown in the first column of the tables. In addition, for each value of M , $\text{var}[u(\mathbf{x}, \boldsymbol{\xi})]$ was estimated at the spatial grid points

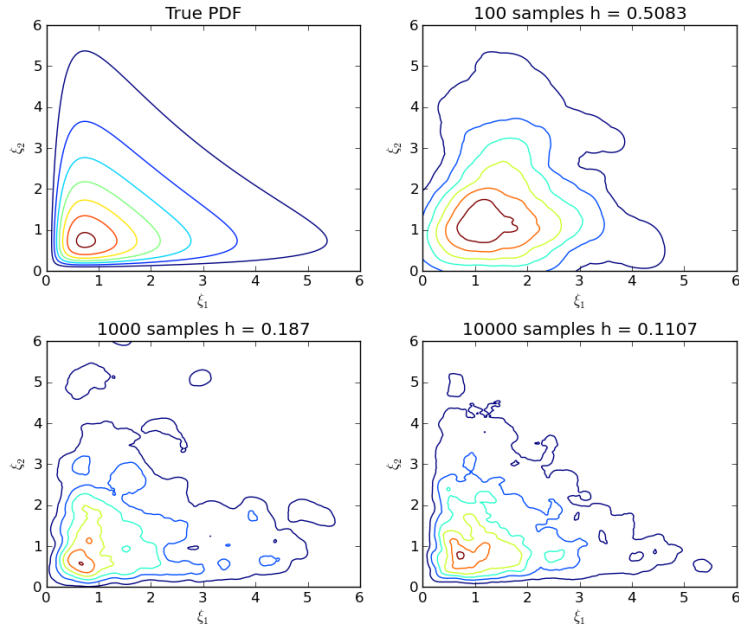


FIG. 5.6. Kernel density estimates for varying numbers of samples.




FIG. 5.7. $\|u(\mathbf{x}, \xi) - A(u(\mathbf{x}, \xi))\rho(\xi)\|_{L^2(D) \times L^\infty(\Gamma)}$ versus the number of collocation points

using 20,000 samples. Equation (4.10) can then be used to compute a 95% confidence bound of the Monte-Carlo error. This estimate is shown in the first column of Tables 5.1, 5.2, and 5.3 beneath the number of samples used to construct the Monte-Carlo

N	τ				
	5×10^{-2}	1×10^{-3}	5×10^{-4}	1×10^{-4}	5×10^{-5}
100 8.43×10^{-2}	5.25×10^{-3} (28)	2.23×10^{-4} (212)	1.18×10^{-4} (301)	9.42×10^{-6} (813)	9.42×10^{-7} (1169)
500 3.78×10^{-2}	5.47×10^{-3} (28)	2.71×10^{-4} (211)	9.84×10^{-5} (315)	1.12×10^{-5} (777)	1.76×10^{-6} (1210)
1000 2.67×10^{-2}	4.29×10^{-3} (33)	2.36×10^{-4} (200)	1.24×10^{-4} (297)	9.78×10^{-6} (762)	2.61×10^{-6} (1207)
5000 1.19×10^{-2}	4.36×10^{-3} (33)	3.88×10^{-4} (172)	1.36×10^{-4} (286)	1.67×10^{-5} (745)	4.73×10^{-6} (1104)
20000 5.96×10^{-3}	4.32×10^{-3} (33)	2.73×10^{-4} (180)	1.30×10^{-4} (294)	1.09×10^{-5} (780)	3.58×10^{-6} (1107)

TABLE 5.1

Monte-Carlo error (left) and $\|\frac{1}{N} \sum_{i=1}^N u(x, \xi^{(i)}) - \mathcal{A}(u)(x, \xi^{(i)})\|_{l^2(D)}$, 4 parameter problem

N	τ				
	5×10^{-2}	1×10^{-3}	5×10^{-4}	1×10^{-4}	5×10^{-5}
100 9.08×10^{-2}	7.66×10^{-3} (76)	8.86×10^{-4} (1026)	4.41×10^{-4} (1655)	4.48×10^{-5} (5026)	8.28×10^{-6} (8111)
500 4.06×10^{-2}	7.13×10^{-3} (92)	6.08×10^{-4} (1170)	3.36×10^{-4} (1189)	2.34×10^{-5} (5773)	1.01×10^{-5} (9404)
1000 2.87×10^{-2}	9.19×10^{-3} (59)	6.03×10^{-4} (1216)	2.65×10^{-4} (1989)	1.95×10^{-5} (5996)	1.77×10^{-5} (9664)
5000 1.28×10^{-2}	7.16×10^{-3} (93)	6.62×10^{-4} (1120)	3.03×10^{-4} (2041)	2.04×10^{-5} (6095)	1.02×10^{-5} (9787)
20000 6.42×10^{-3}	7.25×10^{-3} (93)	6.27×10^{-4} (1187)	2.66×10^{-4} (2127)	1.96×10^{-5} (6050)	5.67×10^{-6} (9942)

TABLE 5.2

Monte-Carlo error (left) and $\|\frac{1}{N} \sum_{i=1}^N u(x, \xi^{(i)}) - \mathcal{A}(u)(x, \xi^{(i)})\|_{l^2(D)}$, 10 parameter problem

N	τ				
	5×10^{-2}	1×10^{-3}	5×10^{-4}	1×10^{-4}	5×10^{-5}
100 9.14×10^{-2}	1.64×10^{-2} (41)	1.65×10^{-3} (878)	2.15×10^{-3} (1299)	5.81×10^{-4} (4126)	2.39×10^{-4} (6958)
500 4.09×10^{-2}	1.45×10^{-2} (41)	2.77×10^{-3} (1045)	1.38×10^{-3} (1738)	3.75×10^{-4} (5545)	1.67×10^{-4} (9106)
1000 2.89×10^{-2}	8.45×10^{-3} (119)	1.46×10^{-3} (1618)	9.02×10^{-4} (2622)	1.66×10^{-4} (8580)	7.13×10^{-5} (14012)
5000 1.29×10^{-2}	8.70×10^{-3} (156)	9.58×10^{-4} (2459)	4.99×10^{-4} (4169)	7.88×10^{-5} (13389)	2.59×10^{-5} (22276)
20000 6.46×10^{-3}	7.25×10^{-3} (193)	6.52×10^{-4} (3108)	3.38×10^{-4} (4991)	3.48×10^{-5} (15963)	2.35×10^{-5} (26081)

TABLE 5.3

Monte-Carlo error (left) and $\|\frac{1}{N} \sum_{i=1}^N u(x, \xi^{(i)}) - \mathcal{A}(u)(x, \xi^{(i)})\|_{l^2(D)}$, 20 parameter problem

estimate.

The other columns of the tables contain results for adaptive KDE collocation where the kernel density estimates are generated using the same set of sample points used for the Monte-Carlo simulation. The total error for this method is bounded by (4.12). The term $\|\epsilon_{MC}\|_{l^2(D)}$ is estimated by the 95% confidence bound in the first column of the tables, as discussed in the previous paragraph. The other quantities in the table are the $l^2(D)$ -norm of the sample mean interpolation error, $\|\epsilon_{interp}\|_{l^2(D)}$, in

the top of each box, together with (in parentheses) the number of collocation points N_{interp} used to construct $\mathcal{A}(u)$. For example, the second from left entry in the bottom row of Table 5.3 shows that for the 20-parameter problem and the 20,000 sample set, $\mathcal{A}(u)$ was constructed using 3,108 collocation points and $\|\epsilon_{interp}\|_{l^2(D)} = 6.52 \times 10^{-4}$.

The costs of the two methods are essentially determined by the number of PDE solves required, N for the Monte-Carlo simulation and N_{interp} for adaptive KDE collocation. In the tables, the number of collocation points N_{interp} in parentheses are shown in bold typeface when they are smaller than the number of samples. For such cases, if $\|\epsilon_{interp}\|_{l^2(D)}$ is significantly smaller than $\|\epsilon_{MC}\|_{l^2(D)}$, then adaptive KDE collocation is less expensive than Monte-Carlo simulation. It can be seen from the results that the savings can be significant when the number of samples increases. For example, the second from left entry in the bottom row of Table 5.3 shows that (by (4.12)) the error in mean for the adaptive collocation method is bounded by $\|\epsilon_{interp}\|_{l^2(D)} + \|\epsilon_{MC}\|_{l^2(D)} = 7.11 \times 10^{-3}$ while only requiring 3,108 PDE solves, an error comparable in magnitude to that obtained with the Monte-Carlo method (6.46×10^{-3}) with 20,000 solves.

We also note that these results suggest that the factor $\log_2(|\Delta\theta^k|)^{3(M-1)}$ in the estimate (4.3) may be pessimistic for many problems of interest. Care must be taken when using the predictor method not to over-resolve the interpolant when one only has access to only a small amount of data. Doing so results in an interpolant that is too accurate given the number of samples available and results in wasted computation. This is the case in the right-hand columns of the tables where the interpolant is being resolved to a much higher level of accuracy than the associated Monte-Carlo error bound.

6. Conclusions. We have presented a new adaptive sparse grid collocation method based on the method proposed in [14] that can be used when the joint PDF of the stochastic parameters is not available and all one has access to is a finite set of samples from that distribution. It is shown that in this case a kernel density estimate can provide a mechanism for driving the refinement of an adaptive sparse grid collocation strategy. Numerical experiments show that in cases involving a large number of samples it can be economical to construct a surrogate to the unknown function using fewer function evaluations and then to perform the Monte-Carlo method on that surrogate.

Acknowledgments. The authors would like to thank Elisabeth Ullmann for her careful reading and helpful comments during the preparation of this work.

REFERENCES

- [1] I. BABUŠKA, F. NOBILE, AND R. TEMPONE, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM Journal on Numerical Analysis, 45 (2007), pp. 1005–1034.
- [2] I. BABUŠKA, R. TEMPONE, AND G. E. ZOURARIS, *Galerkin finite element approximations of stochastic elliptic partial differential equations*, SIAM Journal on Numerical Analysis, 42 (2004), pp. 800–825.
- [3] A. BARTH, C. SCHWAB, AND N. ZOLLINGER, *Multi-level monte carlo finite element method for elliptic pdes with stochastic coefficients*, Numerische Mathematik, 119 (2011), pp. 123–161.
- [4] R. P. BRENT, *Algorithms for Minimization Without Derivatives*, Prentice-Hall, New Jersey, 1973.
- [5] R. E. CAFLISCH, *Monte carlo and quasi-monte carlo methods*, Acta Numerica, 7 (1998), pp. 1–49.

- [6] K. A. CLIFFE, M. B. GILES, R. SCHEICHL, AND A. L. TECKENTRUP, *Multilevel monte carlo methods and applications to elliptic pdes with random coefficients*. Accepted for publication in *Computing and Visualization in Science*, 2011.
- [7] H. C. ELMAN, C. W. MILLER, E. T. PHIPPS, AND R. S. TUMINARO, *Assessment of collocation and galerkin approaches to linear diffusion equations with random data*, *International Journal for Uncertainty Quantifications*, 1 (2011), pp. 19–33.
- [8] R. GHANEM AND P. SPANOS, *Stochastic Finite Elements: A Spectral Approach*, Springer-Verlag, New York, 1991.
- [9] M. GRIGORIU, *Probabilistic models for stochastic partial differential equations*, *Journal of Computational Physics*, 229 (2010), pp. 8406–8429.
- [10] J. HAN AND M. KAMBER, *Data Mining: Concepts and Techniques*, Morgan Kaufmann, San Francisco, 2000.
- [11] A. KLIMKE, *Uncertainty Modeling Using Fuzzy Arithmetic and Sparse Grids*, PhD thesis, Universität Stuttgart, 2006.
- [12] A. KLIMKE AND B. WOHLMUTH, *Algorithm 847: spinterp: piecewise multilinear hierarchical sparse grid interpolation in MATLAB*, *ACM Transactions on Mathematical Software*, 31 (2005), pp. 561–579.
- [13] M. LOEVE, *Probability Theory*, vol. 2, Springer-Verlag, New York, 4th ed., 1978.
- [14] X. MA AND N. ZABARAS, *An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations*, *Journal of Computational Physics*, 228 (2009), pp. 3084–3113.
- [15] ———, *High-dimensional stochastic model representation technique for the solution of stochastic pdes*, *Journal of Computational Physics*, 229 (2010), pp. 3884–3915.
- [16] N. METROPOLIS AND S. ULAM, *The Monte-Carlo method*, *Journal of the American Statistical Association*, 44 (1949), pp. 335–341.
- [17] F. NOBILE AND R. TEMPONE, *A sparse grid stochastic collocation method for partial differential equations with random input data*, *SIAM Journal on Numerical Analysis*, 45 (2008), pp. 2309–2345.
- [18] F. NOBILE, R. TEMPONE, AND C. WEBSTER, *An anisotropic sparse grid collocation algorithm for the solution of stochastic differential equations*, *SIAM Journal on Numerical Analysis*, 46 (2008), pp. 2411–2442.
- [19] B. W. SILVERMAN, *Density Estimation for Statistics and Data Analysis*, Chapman and Hall, London, 1986.
- [20] S. SMOLYAK, *Quadrature and interpolation formulas for tensor products of certain classes of functions*, *Soviet Math Dokl.*, 4 (1963), pp. 240–243.
- [21] X. WAN AND G. E. KARNIADAKIS, *An adaptive multi-element generalized polynomial chaos method for stochastic differential equations*, *Journal of Computational Physics*, 209 (2005), pp. 617–642.
- [22] X. WAN AND G. E. KARNIADAKIS, *Solving elliptic problems with non-Gaussian spatially-dependent random coefficients*, *Computer Methods in Applied Mechanics and Engineering*, 198 (2009), pp. 1985–1995.
- [23] D. XIU AND J. HESTHAVEN, *High-order collocation methods for differential equations with random inputs*, *SIAM Journal on Scientific Computing*, 27 (2005), pp. 1118–1139.