Enhancing ℓ_1 -minimization estimates of polynomial chaos expansions using basis selection

Jakeman, J.D.^{a,*}, Eldred, M.S.^a, Sargsyan, K.^b

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

^b*Sandia National Laboratories, Livermore, CA 94550, United States*

Abstract

In this paper we present a basis selection method that can be used with ℓ_1 -minimization to adaptively determine the large coefficients of polynomial chaos expansions (PCE). The adaptive construction produces anisotropic basis sets that have more terms in important dimensions and limits the number of unimportant terms that increase mutual coherence and thus degrade the performance of ℓ_1 -minimization. The important features and the accuracy of basis selection are demonstrated with a number of numerical examples. Specifically, we show that for a given computational budget, basis selection produces a more accurate PCE than would be obtained if the basis is fixed a priori. We also demonstrate that basis selection can be applied with non-uniform random variables and can leverage gradient information.

Keywords: uncertainty quantification, stochastic collocation, polynomial chaos, ℓ_1 -minimization, sparsity, adaptivity, basis selection

1. Introduction

Quantifying uncertainty in a computational model is essential to building the confidence of stakeholders in the predictions of that model. Sources of uncertainty in model predictions can be broadly grouped into two classes, uncertainty arising from model structure and uncertainty arising from the model parameterization. The effect of these uncertainties must be traced through the model and the effect on the model output (prediction) needs to be quantified. In this paper we will present a method for quantifying parametric uncertainty that utilizes the strengths of Polynomial Chaos Expansions (PCE) and ℓ_1 -minimization.

When the computational cost of a simulation model is large, the most popular and effective means of quantifying parametric uncertainty is to construct an approximation of the response of the model output to variations in the model input. Once built, this surrogate can be interrogated cheaply, without further model evaluations, to obtain statistics of interest such as model output moments and distributions. Within the computational

*Corresponding author

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

Preprint submitted to *Journal of Computational Physics*

November 21, 2014

science community, the most widely adopted approximation methods used for Uncertainty Quantification (UQ) are based on generalized polynomial chaos expansions [25, 46], sparse grid interpolation [27, 31] and Gaussian process models [37].

Polynomial chaos expansions represent a response surface as a linear combination of orthonormal multivariate polynomials. The choice of the orthonormal polynomials is related to the distribution of the model input variables. Provided sufficient smoothness conditions are met, PCEs exhibit high rates of convergence. Indeed in some cases even exponential convergence can be obtained [1, 46]. In this paper we will focus on PCEs as they allow one to leverage the advantages of ℓ_1 -minimization for computing approximations from limited data.

The stochastic Galerkin [25, 46] and stochastic collocation [2, 33, 41, 45] methods are the two main approaches for approximating the PCE coefficients. The former is typically intrusive and so is only feasible when one has the ability to modify the code used to solve the governing equations of the model. Stochastic collocation, however, is a non-intrusive sampling based approach that allows the computational model to be treated as a black box. In this paper we focus on stochastic collocation which involves running the computational model with a set of realizations of the random parameters and constructing an approximation of corresponding model output.

Pseudo-spectral projection [16, 17], sparse grid interpolation [23, 27, 31, 35], probabilistic multi-element methods [22] are stochastic collocation methods which have been used effectively in many situations. These methods, however, all require structured samples and/or the ability to iteratively determine the collocation points.

Recently ℓ_1 -minimization [11, 12, 18, 19] and other sparse approximation methods [14, 42] have been shown to be an effective method for approximating PCE coefficients from small number of and possibly arbitrarily positioned collocation nodes [7, 20, 32, 36, 39, 49]. These methods are very effective when the number of non-zero terms in the PCE approximation of the model output is small (i.e. sparse) or the magnitude of the PCE coefficients decays rapidly (i.e. compressible).

The efficacy of ℓ_1 -minimization when used to estimate PCE coefficients is dependent on the rate of the decay of the PCE coefficients, the characteristics of the stochastic collocation samples and the truncation of the PCE. The decay of the coefficients is dependent on the sparsity of the model response in the chosen basis. In this paper we take the typical approach [7, 20] and adopt the Askey scheme [46], which selects the polynomial basis according to the distribution of the uncertain variables. The Askey scheme is desirable as it facilitates the analytical computation of moments and sensitivity indices from a PCE once it is built [40]. Using the Askey scheme the type of PCE basis is set and cannot be adjusted to enhance ℓ_1 -recovery, however the truncation of the PCE and the sampling of the model inputs can both be controlled.

Recently some attention has been given to designing sampling strategies to increase the accuracy of sparse PCE [38, 47, 48]. Almost no attention, however, has been given to the effect of the PCE truncation when using ℓ_1 -minimization. Typically, when using ℓ_1 -minimization, a total-degree truncation is applied to PCE. However the number of terms in this basis grows factorially with the number of model parameters. This fast growth in the number of basis terms significantly affects the ability of ℓ_1 -minimization to accurately approximate PCE coefficients. To reduce the growth of a PCE basis in high dimensions a hyperbolic cross PCE truncation can be employed [7]. However, despite the slower growth of the hyperbolic truncation it can perform poorly when the ‘true’ PCE

has large coefficients associated with interaction basis terms.

The goal of this paper is to present a basis selection algorithm that adaptively determines a set of PCE basis terms that enable accurate approximation of PCE coefficients using ℓ_1 -minimization. Specifically, we aim to:

- Present an iterative algorithm for selecting a polynomial chaos basis set that, for a given computational budget, produces a more accurate PCE than would be obtained if the basis is fixed a priori.
- Demonstrate numerically that in high dimensions, for which high-order total-degree PCE bases are infeasible, basis selection allows the accurate identification of high-order terms that cannot be captured by a low-order total-degree basis.
- Demonstrate numerically that even for lower dimensional problems, for which high-order total-degree PCE bases are feasible, basis selection still produces more accurate results than a priori fixed basis sets.
- Show that basis selection can leverage function gradients, that for a given computational budget, will produce more accurate approximations than an approximation based solely on function values.
- Illustrate that basis selection can be applied with non-uniform random variables.

The remainder of this paper is organized as follows: Section 2 provides a brief summary of PCEs; Section 3 discusses how to use ℓ_1 -minimization for building a PCE and the need to move away from a priori-fixed PCE truncations in higher dimensions; Section 4 proposes a new method for iteratively defining PCE truncations; the properties and effectiveness of the proposed method are demonstrated numerically in Section 5; and conclusions are presented in Section 6.

2. Polynomial chaos expansions

Polynomial Chaos methods represent both the model inputs $\boldsymbol{\theta} = (\theta_1, \dots, \theta_{\tilde{d}})$ and model output $f(\boldsymbol{\theta})$ as an expansion of orthonormal polynomials of random variables $\boldsymbol{\xi} = (\xi_1, \dots, \xi_d)$. Specifically we represent the random inputs as

$$\theta_n \approx \sum_{i=1}^{N_{\theta_n}} \beta_i \phi_i(\boldsymbol{\xi}), \quad n = 1, \dots, \tilde{d} \quad (1)$$

and the model output as

$$f(\boldsymbol{\theta}(\boldsymbol{\xi})) \approx \hat{f}(\boldsymbol{\xi}) = \sum_{i=1}^N \alpha_i \phi_i(\boldsymbol{\xi}). \quad (2)$$

We refer to (1) and (2) as a polynomial chaos expansion (PCE). The PCE basis functions $\{\phi_i(\boldsymbol{\xi})\}$ are tensor products of orthonormal polynomials which are chosen to be orthonormal with respect to the probability distribution function $\rho(\boldsymbol{\xi})$ of the random vector $\boldsymbol{\xi}$. That is

$$(\phi_i(\boldsymbol{\xi}), \phi_j(\boldsymbol{\xi})) = \int_{I_{\boldsymbol{\xi}}} \phi_i(\boldsymbol{\xi}) \phi_j(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi} = \delta_{ij}$$

where $I_{\boldsymbol{\xi}}$ is the range of the random variables.

The random variable (germ) $\boldsymbol{\xi}$ of the PCE is typically related to the distribution of the input variables. For example, if the one-dimensional input variable θ is uniform on $[a, b]$ then ξ is also chosen to be uniform on $[-1, 1]$ and ϕ are chosen to be Legendre polynomials such that $\theta = \beta_1 + \beta_2\xi = (b + a)/2 + \xi(b - a)/2$. For simplicity and without loss of generality, we will assume that $\boldsymbol{\xi}$ has the same distribution as $\boldsymbol{\theta}$ and thus we can use the two variables interchangeably (up to a linear transformation which we will ignore).

The rate of convergence is dependent on the regularity of the response surface. If $f(\boldsymbol{\xi})$ is analytical with respect to the random variables then (2) converges exponentially in $L_2(\rho(\boldsymbol{\xi}))$ -sense [6].

In practice the PCE (2) must be truncated. The most common approach is to set a degree p and retain only the multivariate polynomials of degree at most p . Rewriting (2) using the typical multi-dimensional index notation

$$f(\boldsymbol{\xi}) \approx \hat{f}(\boldsymbol{\xi}) = \sum_{\boldsymbol{\lambda} \in \Lambda} \alpha_{\boldsymbol{\lambda}} \phi_{\boldsymbol{\lambda}}(\boldsymbol{\xi}) \quad (3)$$

the total-degree basis of degree p is given by

$$\Lambda = \Lambda_{p,q}^d = \{\phi_{\boldsymbol{\lambda}} : \|\boldsymbol{\lambda}\|_q \leq p\}, \quad \boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_d) \quad (4)$$

with $q = 1$. The number of terms in this total-degree basis

$$\text{card } \Lambda_{p,1}^d \equiv P = \binom{d+p}{d}$$

grows factorially with dimension. This rapid growth limits the applicability of the total-degree basis to moderate dimensions or low degree polynomials in higher dimensions.

The authors of [7] propose using hyperbolic index sets, (4) with $q < 1$, to slow the growth of the PCE basis with dimensionality. The use of hyperbolic indices assumes that the contribution to variance from the interaction between the random variables decays rapidly as the number of variables involved in the interaction increases. Figure 1 shows both a three-dimensional total-degree index set and a three-dimensional hyperbolic index set. It is clear that for a given degree p the hyperbolic index set has many less terms than the total-degree polynomial basis. However the smaller basis size requires omitting polynomial terms interaction terms, that is indices $\boldsymbol{\lambda}$ with at least two $\lambda_n > 0, n = 1, \dots, d$. When a function has large non-zero PCE coefficients corresponding to these missing multivariate basis terms, the hyperbolic index set may be an inappropriate form of truncation. Ideally the basis set should be adapted to the function being approximated.

3. ℓ_1 -minimization

The coefficients of a polynomial chaos expansion can be approximated effectively using ℓ_1 -minimization. Specifically, given a small set of M unstructured realizations $\boldsymbol{\Xi} = \{\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_M\}$, with corresponding model outputs $\mathbf{f} = (f(\boldsymbol{\xi}_1), \dots, f(\boldsymbol{\xi}_M))^T$, we would like to find a solution that satisfies

$$\Phi \boldsymbol{\alpha} \approx \mathbf{f}$$

Figure 1: (Left) A total-degree index set $\Lambda_{6,1}^3$. (Right) A hyperbolic index set $\Lambda_{6,1/2}^3$. Each cube represents a 3-dimensional index $\lambda = (\lambda_1, \lambda_2, \lambda_3)$ in Λ . The value of each λ is given on the top of each cube.

where $\alpha = (\alpha_{\lambda_1}, \dots, \alpha_{\lambda_N})^T$ denotes the vector of PCE coefficients and Φ denotes the Vandermonde matrix with entries $\Phi_{ij} = \phi_j(\xi_i)$, $i = 1, \dots, M$, $j = 1, \dots, N$.

When the model $f(\xi)$ is high-dimensional and computationally expensive, and non-adaptive basis truncation rules are employed, the number of model simulations that can be generated is much smaller than the number of unknown PCE coefficients, i.e $M \ll N$. Under these conditions, finding the PCE coefficients is ill-posed and we must impose some form of regularization to obtain a unique solution. One way to regularize is to enforce sparsity in a PCE.

A polynomial chaos expansion is defined as s -sparse when $\|\alpha\|_0 \leq s$, i.e. the number of non-zero coefficients does not exceed s . ℓ_1 -minimization provides a means of identifying *sparse* coefficient vectors from a limited amount of simulation data. In practice, not many simulation models will be truly sparse, but PCE are often *compressible*, that is the magnitude of the coefficients decays rapidly or alternatively most of the PCE variance is concentrated in a few terms. Compressible vectors are well represented by sparse vectors and thus the coefficients of compressible PCE can also be recovered accurately using ℓ_1 -minimization.

ℓ_1 -minimization attempts to find the dominant PCE coefficients by solving the following optimization problem

$$\alpha^* = \arg \min_{\alpha} \|\alpha\|_1 \quad \text{such that} \quad \|\Phi\alpha - \mathbf{f}\|_2 \leq \varepsilon \quad (5)$$

This ℓ_1 -minimization problem is often referred to as Basis Pursuit Denoising. The problem obtained by setting $\varepsilon = 0$, to enforce interpolation, is termed Basis Pursuit. There is a close connection between (5) and Least Absolute Shrinkage Operator (LASSO) [42] well known in the statistics literature. Indeed these problems are equivalent under certain conditions [19].

3.1. ℓ_1 -minimization algorithms

Numerous algorithms [5, 7, 13, 34, 43] exist for solving (5) which are all stable and accurate under certain well defined conditions. In this paper we will use the greedy algorithm Orthogonal Matching Pursuit (OMP) [14] to estimate PCE coefficients. We remark, however, that the basis selection procedure presented in this paper can be used in conjunction with most ℓ_1 -minimization algorithms.

OMP requires stronger theoretical conditions than some of its counterparts [10]. For example, for a sufficiently small ratio of the number of samples to sparsity, basis pursuit can guarantee recovery of all s -sparse polynomials with high probability, whereas, for a fixed set of samples Ξ , OMP can guarantee recovery of at least one sparse polynomial but not all [29]. However despite this theoretical difference, in practice OMP can still obtain comparable accuracy to non-approximate algorithms such as BP [29]. Moreover OMP has a much faster execution speed than most algorithms, which makes OMP more amenable to cross validation which can be used to estimate optimal method parameters such as the tolerance ε of (5) as discussed in the following section.

3.1.1. Hyper-parameter estimation via cross validation

Accurately computing the coefficients of a polynomial chaos expansion requires determining a ‘good’ truncation set Λ and specifying the tolerance ε in the Basis Pursuit DeNoising problem (5). Cross validation has been shown to be effective at aiding these choices. Specifically cross validation has been used in the past to estimate the polynomial degree p of a hyperbolic expansion [7] and to estimate the tolerance ε of (5) [7, 8, 20, 32, 44]¹.

In this paper we will use $K = 10$ fold cross validation to choose the values of sets of hyper-parameters γ . The number and type of hyper-parameters is dependent on the ℓ_1 -minimization method used in conjunction with cross validation. As an example consider solving in (5) using an a priori fixed total-degree basis $\Lambda_{p,1}^d$. The hyper-parameters that can be estimated using cross validation are the degree p and the tolerance ε , that is $\gamma = (p, \varepsilon)$.

Let $\zeta : \{1, \dots, M\} \rightarrow \{1, \dots, K\}$ be an indexing function that determines the partition of the training data. Furthermore let $\hat{f}^{-\zeta}$ be the PCE approximation built on the data with the ζ part removed, then the cross validation error is given by

$$e_{cv}(\gamma) = \frac{1}{M} \sum_{k=1}^K e_{\zeta(k)}, \quad e_{\zeta(k)} = \sum_{j \in \zeta(k)} (y_j - \hat{f}^{-\zeta(k)}(x_j))^2 \quad (6)$$

To compute e_{cv} we divide the data pairs $(\Xi, \mathbf{f})$, based upon the randomly chosen partitions $\zeta(k)$, into K sets (folds) of equal size $(\Xi_k, \mathbf{f}_k)$, $k = 1, \dots, K$. A PCE $\hat{f}^{-\zeta(k)}$, is then built on the training data $\Xi_t = \Xi \setminus \Xi_k$ with the k -th fold removed, using the hyper-parameters γ . The remaining data $\Xi_v = \Xi_k$ is then used to estimate the prediction error. To estimate the hyper-parameters γ we search over a set of possible values for γ and select $\gamma = \arg \min_{\gamma} e_{cv}(\gamma)$.

¹The choice of ε can significantly affect the accuracy of the PCE obtained using (5). Reducing ε can lead to over-fitting, whilst higher values of ε can deteriorate the accuracy of the approximation.

Figure 2 presents a typical example illustrating the change in the ℓ_2 error of a PCE with fixed degree, as the tolerance ε is decreased. The figure also plots the cross validation error which is a good indicator of the ℓ_2 error behavior. The vertical line represents the tolerance chosen by cross validation and the horizontal line is the ε_{ℓ_2} error in the resulting PCE. The result shown is typical. There is a bias (underestimation of ε_{ℓ_2}) in the cross validation estimate, yet despite this bias cross validation consistently chooses a tolerance that produces a near minimal error.

Figure 2: The use of cross validation to select the truncation tolerance ε for (11) with coefficients $c^{(1)}$. The vertical line represents the tolerance chosen by cross validation and the horizontal line is the ε_{ℓ_2} error in the resulting PCE. The coefficients of the 3rd order total-degree basis were obtained using $M = 200$ uniform samples. Accuracy is measured using the ℓ_2 norm computed at 100,000 Latin-hypercube samples (see Section 5).

We remark that 10-fold cross validation can over-estimate the true approximation error, however we found that this was not a significant issue for the functions we tested, which is consistent with other studies [9, 28]. One could use a larger number of folds, for example $K = M$ (leave-one-out cross validation), however any reduction in bias will likely be accompanied by greater variance [26]. Moreover setting K much larger than 10 significantly increases the computational effort required to perform cross validation.

3.2. Recoverability of ℓ_1 -minimization

The ability of ℓ_1 -minimization to accurately determine the large coefficients of the PCE is determined by the properties of the matrix Φ and the sparsity of PCE representation of the model response $f(\xi)$. The sparsity of the model response is dependent on the PCE basis. In this paper we take the typical approach [7, 20] and adopt the Askey scheme [46], which fixes the polynomial basis based on the distribution of the uncertain variables, thereby facilitating analytical computation of moments from the PCE and sensitivity analysis [40]. Using the Askey scheme the PCE basis and thus the sparsity of the

model response is fixed, however the properties of the Φ are influenced by the selection of the realizations $\{\xi_i\}_{i=1}^M$ and the truncation Λ .

Mutual coherence is one measure often used to indicate the ability of ℓ_1 -minimization to find a sparse solution. The mutual coherence of a matrix $\Phi \in \mathbb{R}^{M \times N}$ with columns $\tilde{\phi}_j$ is

$$\mu(\Phi) = \max_{1 \leq j < k \leq N} \frac{|\tilde{\phi}_j^T \tilde{\phi}_k|}{\|\tilde{\phi}_j\|_2 \|\tilde{\phi}_k\|_2} \quad (7)$$

and is a measure of the maximum correlation between any two columns in the matrix. ℓ_1 -minimization will obtain a better estimate of the PCE coefficients if the mutual coherence of Φ is small. Intuitively, if two columns are closely correlated the mutual coherence will be large and it will be impossible, in general, to distinguish whether the energy in the signal comes from one or the other.

The *restricted isometry property* (RIP) [13], quantified by the restricted isometry constant δ is another measure of the recoverability of the matrix Φ . For each $s = 1, 2, \dots$ the isometry constant δ_s of a matrix Φ is the smallest number such that

$$(1 - \delta_s) \|\alpha_s\|_2^2 \leq \|\Phi \alpha_s\|_2^2 \leq (1 + \delta_s) \|\alpha_s\|_2^2 \quad (8)$$

for all vectors α_s with s non-zero entries. This is equivalent to requiring that the eigenvalues of all Grammian matrices $\Phi_{\Lambda_s}^T \Phi_{\Lambda_s}$ lie between $[1 - \delta_s, 1 + \delta_s]$, where Φ_{Λ_s} are $M \times s$ submatrices of Φ . The restricted isometry property measures the ability of Φ to preserve the lengths of s -sparse vectors. The RIP can be intuitively thought of as a measure of s -wise coherence as opposed to mutual coherence which is a measure of pair-wise coherence.

3.2.1. PCE truncation

Naively choosing a large degree p can cause a degradation in the accuracy of the PCE coefficients. Figure 3 demonstrates that both the mutual coherence and the 10-sparse RIP constant δ_{10} of the Vandermonde matrix Φ increases as the number of basis terms P increases.²

Increases in mutual coherence and RIP constant are correlated with a decrease in the accuracy of PCE coefficients recovered by ℓ_1 -minimization. Figure 4 demonstrates that, for a fixed number of samples, as the number of terms N and, consequently, the mutual coherence and RIP constant increase (see Figure 3), the PCE recovered by ℓ_1 -minimization becomes less accurate. Figure 4 also illustrates that the accuracy of the PCE depends upon the degree of the basis used. As the number of samples M increases, ℓ_1 -minimization is able to recover more dominant coefficients and a higher degree should be used. However for a given number of samples increasing the degree does not always lead to a reduction in error. For example when $M = 60$ the PCE of total-degree $p = 4$ has the smallest error and at $M = 120$ setting $p = 7$ produces the smallest error. It is not until $M = 240$ that the highest degree basis $p = 8$ produces the smallest error. These results are consistent with the theoretical results in [20] that assert that the number of samples M needed to recover a Legendre PCE of a certain sparsity s increases with the number of terms in the PCE basis.

²The RIP constant reported here is a lower bound found by computing the eigenvalues of 10,000 randomly selected submatrices Φ_{Λ_s} .

Figure 3: Growth of mutual coherence and RIP constant δ_{10} with the number of terms P in a 6-dimensional total-degree multivariate Legendre PCE basis.

3.2.2. Sampling strategies and pre-conditioning

The sampling strategy used to choose the samples Ξ affects the mutual coherence and RIP of Φ and thus can impact the accuracy of the recovered polynomial chaos expansion. To date, the best sampling strategies for ℓ_1 -minimization are random [38, 47]. The nature of the random samples is dependent on the distribution of the random variables ξ , the number of random dimensions, and the degree of the PCE.

The accuracy of ℓ_1 -minimization solutions of (5) can also be improved by the use of pre-conditioning. The pre-conditioned ℓ_1 -minimization problem is given by

$$\alpha = \arg \min_{\alpha} \|\alpha\|_1 \quad \text{such that} \quad \|\mathbf{W}\Phi\alpha - \mathbf{W}\mathbf{f}\|_2 \leq \varepsilon \quad (9)$$

where $\mathbf{W} \in \mathbb{R}^{M \times M}$ is a diagonal matrix with entries chosen to enhance the recovery properties of ℓ_1 -minimization. When recovering s -sparse one-dimensional Legendre polynomials, randomly sampling $\Xi = \{\xi_m\}_{m=1}^M$ from the Chebyshev measure and choosing weights $w_{m,m} = (\pi/2)^{d/2}(1 - \xi_m^2)^{1/4}$, can result in significant increases in the accuracy of the coefficients recovered by ℓ_1 -minimization [38]. In the multivariate setting, however, the benefit of pre-conditioning is less clear [48]. In this paper all numerical results presented are generated without pre-conditioning.

4. Iterative basis selection

When the coefficients of a PCE can be well approximated by a sparse vector, ℓ_1 -minimization is extremely effective at recovering the coefficients of that PCE. It is possible, however, to further increase the efficacy of ℓ_1 -minimization by leveraging real-

Figure 4: The dependence of PCE accuracy on the polynomial degree p . The *degree* of the most accurate expansion is dependent on the number of LHS samples M used to construct the PCE. Results were obtained using orthogonal matching pursuit with cross validation to choose ε applied to the 6-dimensional random oscillator (12). Accuracy is measured using the ℓ_2 norm computed at 100,000 Latin-hypercube samples (see Section 5). The error reported is the average over 20 trials for each design size.

istic models of structural dependencies between the values and locations of the PCE coefficients. For example [3, 21, 30] have successfully increased the performance of ℓ_1 -minimization when recovering wavelet coefficients that exhibit a tree-like structure. In this vein, we propose an algorithm for identifying the large coefficients of PC expansions that form a semi-connected subtree of the PCE coefficient tree.

The coefficients of polynomial chaos expansions often form a multi-dimensional tree. Given an ancestor basis term ϕ_{λ} of degree $\|\lambda\|_1$ we define the indices of its children as $\lambda + \mathbf{e}_k$, $k = 1, \dots, d$, where $\mathbf{e}_k = (0, \dots, 1, \dots, 0)$ is the unit vector co-directional with the k -th dimension. An example of a typical PCE tree is depicted in Figure 5. In this figure, as often in practice, the magnitude of the ancestors of a PCE coefficient is a reasonable indicator of the size of the child coefficient. In practice, some branches (connections) between levels of the tree may be missing. We refer to trees with missing branches as semi-connected trees.

In the following we present a method for estimating PCE coefficients that leverages the tree structure of PCE coefficients to increase the accuracy of coefficient estimates obtained by ℓ_1 -minimization.

4.1. Algorithm

Typically ℓ_1 -minimization is applied to an a priori chosen and fixed basis set Λ . However the accuracy of coefficients obtained by ℓ_1 -minimization can be increased by adaptively selecting the PCE basis.

Figure 5: Tree structure of the coefficients of a two dimensional PCE with $\Lambda_{3,1}^2$. For clarity we only depict one connection per node, but in d dimensions a node of a given degree p will be a child of up to d nodes of degree $p - 1$. For example, not only is the basis $\phi_{[1,1]}$ a child of $\phi_{[1,0]}$ (as depicted) but it is also a child of $\phi_{[0,1]}$

To choose a basis for ℓ_1 -minimization we employ a four step iterative procedure involving restriction, expansion, identification and selection. Pseudo code for the iterative basis selection procedure is outlined in Algorithm 1 and a graphical version of the algorithm is depicted in Figure 6. To emphasize the four stages of the algorithm, each stage is highlighted using the same colors in both Algorithm 1 and Figure 6.

To initiate the basis selection algorithm, we first define a basis set $\Lambda^{(0)}$ and use ℓ_1 -minimization to identify the largest coefficients $\alpha^{(0)}$. The choice of $\Lambda^{(0)}$ can sometimes affect the performance of the basis selection algorithm. For example, if $\Lambda^{(0)}$ contains insufficient terms, for a given sample size M , then the initial basis will not be able to adequately inform the iterative procedure. Conversely, if the cardinality of the initial basis is too large, the increased degradation of the properties of the matrix Φ , such as mutual coherence (see Section 3.2), will result in poor recovery of the initial PCE coefficients which will limit the effectiveness of the basis selection procedure. We found a good choice to be $\Lambda^{(0)} = \Lambda_{p,1}$, where p is the degree that gives $|\Lambda_{p,1}^d|$ closest to $10M$, i.e. $\Lambda_{p,1}^d = \arg \min_{\Lambda_{p,1}^d \in \{\Lambda_{1,1}^d, \Lambda_{2,1}^d, \dots\}} ||\Lambda_{p,1}^d| - 10M|$.

Given a basis $\Lambda^{(k)}$ and corresponding coefficients $\alpha^{(k)}$ we reduce the basis to a set $\Lambda_\varepsilon^{(k)}$ containing only the terms with non-zero coefficients. This restricted basis is then expanded T times using an algorithm which we will describe in Section 4.1.1. ℓ_1 -minimization is then applied to each of the expanded basis sets $\Lambda^{(k,t)}$ for $t = 1, \dots, T$. Each time ℓ_1 -minimization is used, we employ cross validation to choose ε . Therefore, at every basis set considered during the evolution of the algorithm we have a measure of the expected accuracy of the PCE coefficients. At each step in the algorithm the basis set that results in the lowest cross validation error is chosen. Typically the algorithm

Figure 6: Graphical depiction of the basis adaptation algorithm.

will choose the basis set that corresponds to $T = 1$ or $T = 2$.

Algorithm 1: $\Lambda^*, \alpha^* = \text{BASIS_SELECTION}[\Phi, \mathbf{f}, \varepsilon]$

$\Lambda^* = \Lambda^{(0)} = \Lambda_{p,1}^d = \arg \min_{\Lambda_{p,1}^d \in \{\Lambda_{1,1}^d, \Lambda_{2,1}^d, \dots\}} \left| |\Lambda_{p,1}^d| - 10M \right|$

$\alpha^{(0)}, e_{\text{cv}}^{(0)} = \ell_1\text{-minimization}[\Phi(\Lambda^{(0)}), \mathbf{f}]$

$T = 3, e_{\text{cv}}^* = \infty, k = 1$

while *TRUE* **do**

$e_{\text{cv}}^{(k)} = \infty$

$\Lambda^{(k,0)} = \{\lambda : \lambda \in \Lambda^{(k-1)}, \alpha_\lambda^{(k)} \neq 0\}$

for $t \in \{1, \dots, T\}$ **do**

$\Lambda^{(k,t)} = \text{EXPAND}[\Lambda^{(k,t-1)}]$

$\alpha^{(k,t)}, e_{\text{cv}}^{(k,t)} = \ell_1\text{-minimization}[\Phi(\Lambda^{(k,t)}), \mathbf{f}]$

if $e_{\text{cv}}^{(k,t)} < e_{\text{cv}}^{(k)}$ **then**

$e_{\text{cv}}^{(k)} = e_{\text{cv}}^{(k,t)}, \alpha^{(k)} = \alpha^{(k,t)}, \Lambda^{(k)} = \Lambda^{(k,t)}$

end

end

if $e_{\text{cv}}^{(k)} > e_{\text{cv}}^*$ **then**

 TERMINATE

end

$\alpha^* = \alpha^{(k)}, \Lambda^* = \Lambda^{(k)}, e_{\text{cv}}^* = e_{\text{cv}}^{(k)}$

end

4.1.1. Basis expansion

Define $\{\lambda + \mathbf{e}_j : 1 \leq j \leq d\}$ the forward neighborhood of an index λ and similarly let $\{\lambda - \mathbf{e}_j : 1 \leq j \leq d\}$ denote the backward neighborhood. To expand a basis set Λ we must first find the forward neighbors $\mathcal{F} = \{\lambda + \mathbf{e}_j : \lambda \in \Lambda, 1 \leq j \leq d\}$ of all indices $\lambda \in \Lambda$. The expanded basis is then given by

$$\Lambda^+ = \Lambda \cup \mathcal{A}, \quad \mathcal{A} = \{\lambda : \lambda \in \mathcal{F}, \lambda - \mathbf{e}_n \in \Lambda \text{ for } 1 \leq n \leq d, \lambda_k > 1\}$$

where we have used the following admissibility criteria

$$\boldsymbol{\lambda} - \mathbf{e}_n \in \Lambda \text{ for } 1 \leq n \leq d, \lambda_k > 1 \quad (10)$$

to target PCE basis indices that are likely to have large PCE coefficients. A forward neighbor is admissible only if its backward neighbors exist in all dimensions. If the backward neighbors do not exist then ℓ_1 -minimization has previously identified that the coefficients of these backward neighbors are negligible.

The admissibility criterion is explained graphically in Figure 7. In the left graphic, both children of the current index are admissible, because its backwards neighbors exist in every dimension. In the right graphic only the child in the vertical dimension is admissible, as not all parents of the horizontal child exist.

Figure 7: Identification of the admissible indices of an index (red). The indices of the current basis Λ are gray and admissible indices are striped. A index is admissible only if its backwards neighbors exists in every dimension.

At the k -th iteration of Algorithm 1, ℓ_1 -minimization is applied to $\Lambda^{(k-1)}$ and used to identify the significant coefficients of the PCE and their corresponding basis terms $\Lambda^{(k,0)}$. The set of non-zero coefficients $\Lambda^{(k,0)}$ identified by ℓ_1 -minimization is then expanded. The **EXPAND** routine expands an index set by one polynomial degree, but sometimes it may be necessary to expand the basis $\Lambda^{(k)}$ more than once.³ To generate these higher degree index sets **EXPAND** is applied recursively to $\Lambda^{(k,0)}$ up to a fixed number of T times. Specifically, the following sets are generated

$$\Lambda^{(k,t)} = \Lambda^{(k,t-1)} \cup \{\boldsymbol{\lambda} : \boldsymbol{\lambda} - \mathbf{e}_n \in \Lambda^{(k,t-1)}, 1 \leq n \leq d, \lambda_n > 1\}.$$

³The choice of $T > 1$ enables the basis selection algorithm to be applied to semi-connected tree structures as well as fully connected trees. Setting $T > 1$ allows us to prevent premature termination of the algorithm if most of the coefficients of the children of the current set $\Lambda^{(k)}$ are small but the coefficients of the children's children are not.

As the number of expansion steps T increases the number of terms in the expanded basis increases rapidly and degradation in the performance of ℓ_1 -minimization can result (this is similar to what happens when increasing the degree of a total-degree basis). To avoid degradation of the solution, we use cross validation to choose the number of inner expansion steps $t \in [1, T]$.

5. Numerical examples

In this section we use several numerical tests to demonstrate the benefit of the basis selection method. In each example we seek a PCE approximation to a model output given a set of uncertain parameters with a known range or distribution.

We compare the approximations constructed using basis selection against those constructed using a non-adaptive strategy. The non-adaptive strategy consists of generating basis sets $\Lambda_{2,1}^d, \dots, \Lambda_{p,1}^d$ where p is the degree that produces the basis set with a cardinality closest to 100000. ℓ_1 -minimization is then applied to this basis with a cross validation tolerance search to compute the non-zero polynomial coefficients. The resulting basis with the lowest cross validation error is chosen to be the final approximation.

We also compare the non-adaptive and basis selection methods against OMP using a basis oracle. Given an exact PCE representation of a scalar function and M function samples, the oracle basis consists of $N = M$ basis terms corresponding to the N largest coefficients of the exact PCE. For the examples that follow, the exact PCE are not known analytical and so they are computed to machine precision using a dimension adaptive sparse grid. The oracle basis, defined here, will be close to optimal and therefore will serve as a good estimate of the maximum accuracy that can be gained from the use of basis selection.

In order to construct a PCE approximation, the sample design has to be specified. By a design we mean the choice of sample size, M , and the selection $\Xi = \{\xi_i\}_{i=1}^M$. Here we opt for uniform random samples of size M . For small sample sizes M the selection of Ξ significantly affects the performance of any approximation method. Therefore, for each M , twenty different designs are used to build a PCE and the median first and third quartiles and maximum and minimum of the resulting errors are reported using box and whisker plots. We did investigate the utility of using samples drawn from the Chebyshev measure but found that there was no consistent benefit. Even in the cases for which a benefit was observed, the improvement was small relative to the benefit gained from using basis adaptation. This finding is consistent with [48].

To measure the performance of an approximation, we will use the ℓ_2 error (RMSE). Specifically given a set of $Q = 100,000$ Latin-hypercube samples $\Xi_{\text{test}} = \{\xi^{(i)}\}_{i=1}^Q \in I_{\xi}$ and samples of the true function $f(\xi^{(i)})$ and the PCE approximation $\hat{f}(\xi^{(i)})$ we compute

$$\varepsilon_{\ell_2} = \left(\frac{1}{Q} \sum_{i=1}^Q |\hat{f}(\xi^{(i)}) - f(\xi^{(i)})|^2 \right)^{1/2}$$

Note in all examples presented using Legendre polynomials we transform each d dimensional parameter domain I_{ξ} , build points Ξ and test points Ξ_{test} to $[-1, 1]^d$.

5.1. Algebraic test function

Consider the algebraic corner-peak test function [24]

$$f_{\text{CP}}(\mathbf{x}) = \left(1 + \sum_{k=1}^d c_k \xi_k\right)^{-(d+1)}, \quad \boldsymbol{\xi} \in [0, 1]^d \quad (11)$$

This function provides a flexible test that can be used to identify the strengths of the proposed algorithm. Specifically, the coefficients c_k can be used to control the effective dimensionality and the compressibility of these functions. Here we will examine performance using three different choices of $\mathbf{c} = (c_1, \dots, c_d)^T$, specifically

$$c_k^{(1)} = \frac{k - \frac{1}{2}}{d}, \quad c_k^{(2)} = \frac{1}{k^2} \quad \text{and} \quad c_k^{(3)} = \exp\left(\frac{k \log(10^{-8})}{d}\right), \quad k = 1, \dots, d$$

normalizing such that $\sum_{k=1}^d c_k = 0.25$. The coefficients $\mathbf{c}^{(1)}$, $\mathbf{c}^{(2)}$ and $\mathbf{c}^{(3)}$ represent increasing levels of anisotropy and decreasing effective dimensionality. Anisotropy refers to the dependence of the function variability, often measured through variance, on individual parameter dimensions ξ_n . When a function is strongly anisotropic, the majority of the function variance can be attributed to a small set of dimensions. The size of this subset is referred to as the effective dimension.

The performance of the basis selection method is dependent on the compressibility of the model response in the chosen basis. Figure 8 plots the ε_{ℓ_2} error in the polynomial approximations for increasing number of model evaluations. For all three levels of anisotropy the adaptive method produces an expansion no worse than the non-adaptive method, for the same sample size. When anisotropy is introduced the accuracy of basis selection increases relative to the non-adaptive method. The stronger the anisotropy the better the relative performance. Figure 8 also plots the PCE obtained using OMP with an oracle basis. When very weak anisotropy is present $\mathbf{c}^{(1)}$, there is little that can be gained by using a well chosen basis, as evident by the lack of separation between the three convergence curves. However as the strength of the anisotropy is increased the effect of the oracle basis on accuracy becomes much more apparent. When strong anisotropy, $\mathbf{c}^{(3)}$, is present, basis selection is able to obtain the same accuracy as the oracle without a priori information on the truncation of the basis which is required by the oracle. In the moderately anisotropic $\mathbf{c}^{(2)}$ case basis selection does not perform as well as the oracle but does still perform better than the non-adaptive method.

To understand the correlation between anisotropy and the performance of the basis selection method we must consider the structure of the PCE coefficients induced by varying $\mathbf{c}$. Figure 8 (d) plots the decay of the PCE coefficients when sorted by magnitude. As anisotropy increases, so does the rate of decay of the sorted PCE coefficients.

It is the strength of decay that controls the performance of basis selection. When the rate of decay is high then the function is more compressible and thus better suited to being approximated using ℓ_1 -minimization. Anisotropy will often result in compressible coefficients, but it is conceptually possible for models to be compressible without being anisotropic. For some problems such as the elliptic Poisson equation the decay of the PCE coefficients can be calculated a-priori [1, 4, 15] but unfortunately in practice, the decay of the PCE coefficients of a model cannot be determined ahead of time. A practical

(a)

(b)

(c)

(d)

Figure 8: Convergence of the RMSE, with respect to increasing design size M , in the Legendre PCE approximation of the Genz corner-peak function (11) for the three coefficient regimes: (a) $c^{(1)}$, (b) $c^{(2)}$, (c) $c^{(3)}$. The error bars (whiskers) represent the minimum and maximum error over the 20 trials for each design size M and the boxes represent the 1st, 2nd (median) and 3rd quartiles. (d) PCE coefficients, sorted by magnitude, for the three coefficient regimes.

means of identifying the coefficient decay regime would be very useful but is beyond the scope of this paper.

Not only does the rate of coefficient decay affect performance, but so to does the ability of the PCE basis Λ to represent the target function. For example if the ‘true’ PCE has large high degree terms with large coefficients but the basis Λ does not have these high degree terms then the PCE obtained using Λ will not be as accurate as a PCE obtained using a basis that included the important high degree terms.

Figure 9 plots the exact PCE coefficients of the algebraic test function using $\mathbf{c}^{(1)}$ and $\mathbf{c}^{(3)}$. The ‘exact’ were coefficients obtained using a dimension-adaptive sparse grid with 100,000 samples, which resulted in an approximation error below 10^{-8} . In Figure 9 (a), the random variables contribute similarly to the total variance of f_{CP} and so the dominant coefficients are concentrated in the lower degree terms of the PCE, thus a total-degree basis set will perform as well as any alternative. In comparison, the importance of the dimensions of the function shown in Figure 9 (b) decay exponentially with dimension, which results in higher-degree terms with large coefficients in some dimensions. In this coefficient regime, if ℓ_1 -minimization can only be applied with a low degree polynomial (which is true when using a total-degree basis), the accuracy of the resulting PCE, for a given number of samples M , will not be as high as a PCE constructed using a basis that includes the dominant high-degree terms. Basis selection will typically allow identification and recovery of more high-degree coefficients than would be possible if using a total-degree basis.

Figure 9: Comparison of PCE coefficients of $\mathbf{c}^{(1)}$ (a) and $\mathbf{c}^{(3)}$ (b). The black squares represent the ‘exact’ coefficients, the red circles the non-zero coefficients of the basis selection method and the blue diamonds are the non-zero coefficients recovered by the non-adaptive approach. The indices on the x axis are sorted lexicographically by degree. The dashed vertical lines separate the PCE terms into degrees. The horizontal lines represent the ℓ_2 error in the basis selection and non-adapted PCE. The identification of more terms accurately with basis selection results in a smaller error.

5.2. Random oscillator

In this section we investigate the performance of basis selection to quantify uncertainty in a damped linear oscillator subject to external forcing with six unknown parameters. That is,

$$\frac{d^2x}{dt^2}(t, \boldsymbol{\xi}) + \gamma \frac{dx}{dt} + kx = f \cos(\omega t), \quad (12)$$

subject to the initial conditions

$$x(0) = x_0, \quad \dot{x}(0) = x_1, \quad (13)$$

where we assume the damping coefficient γ , spring constant k , forcing amplitude f and frequency ω , and the initial conditions x_0 and x_1 are all uncertain. We solve (12) analytically to avoid consideration of discretization errors in our study.

Defining $\boldsymbol{\xi} = (\gamma, k, f, \omega, x_0, x_1)$ let $\xi_1 \in [0.08, 0.12]$, $\xi_2 \in [0.03, 0.04]$, $\xi_3 \in [0.08, 0.12]$, $\xi_4 \in [0.8, 1.2]$, $\xi_5 \in [0.45, 0.55]$, $\xi_6 \in [-0.05, 0.05]$. For any parameter realization in $I_{\boldsymbol{\xi}}$ the harmonic oscillator will be underdamped. In the following, we choose our quantity of interest to be the position $x(t)$ of the oscillator at $t = 20$ seconds.

Figure 10 (a) depicts the error in the Legendre PCE for increasing design sizes M . The basis selection method clearly outperforms the non-adaptive approach and produces comparable results to the oracle. Again the improvement in performance is associated with a rapid decay of the exact PCE coefficients (see Figure 10).

5.3. Diffusion equation

In this section, we consider the heterogeneous diffusion equation in one-spatial dimension subject to uncertainty in the diffusivity coefficient. This problem has been used as a benchmark in other works [20, 49]. Attention is restricted to one-dimensional physical space to avoid unnecessary complexity. The procedure described here can easily be extended to higher physical dimensions. Consider the following problem with $d \geq 1$ random dimensions:

$$-\frac{d}{dx} \left[a(x, \boldsymbol{\xi}) \frac{du}{dx}(x, \boldsymbol{\xi}) \right] = 1, \quad (x, \boldsymbol{\xi}) \in (0, 1) \times I_{\boldsymbol{\xi}} \quad (14)$$

subject to the physical boundary conditions

$$u(0, \boldsymbol{\xi}) = 0, \quad u(1, \boldsymbol{\xi}) = 0. \quad (15)$$

Furthermore, assume that the random diffusivity satisfies

$$a(x, \boldsymbol{\xi}) = \bar{a} + \sigma_a \sum_{k=1}^d \sqrt{\lambda_k} \phi_k(x) \xi_k, \quad (16)$$

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

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

(a)

(b)

(c)

Figure 10: Convergence of the RMSE, with respect to increasing design size M , in the Legendre PCE approximation of (a) the harmonic oscillator (12) and (b) the solution $u(1/3, \xi)$ of the diffusion equation (14). The error bars (whiskers) represent the minimum and maximum error over the 20 trials for each design size M and the boxes represent the 1st, 2nd (median) and 3rd quartiles. (c) PCE coefficients, sorted by magnitude, for the harmonic oscillator, the diffusion equation and the resistor network.

Figure 11: Resistor network comprised of $d = 2P$ resistances R_i , $i = 1, \dots, d$ of uncertain ohmage and the network is driven by a voltage source providing a known potential V_0 . We are interested in determining voltage at V .

The variability of the diffusivity field (16) is controlled by σ_a and the correlation length l_c which determines the decay of the eigenvalues λ_k . Here we approximate the solution $u(1/3, \xi)$ with $\bar{a} = 0.1$, $d = 14$, $\sigma_a = 0.03$, $l_c = 1/5$, while the uncertain inputs $\xi_k \in [-1, 1]$, $k = 1, \dots, d$ are independent and uniformly distributed random variables. We solve the model (14) using quadratic finite elements with a high enough spatial resolution to neglect discretization errors in our analysis.

Figure 10 (b) plots the error in Legendre PCE approximations built using increasing design sizes M . There is negligible difference between basis selection and the non-adaptive strategy, but there is also negligible difference between these methods and the oracle, indicating there is not much improvement that can in principle be gained from basis selection. The negligible improvement is due to the fact that the ‘exact’ PCE is not very compressible, as can be seen from Figure 10. The lack of compressibility means that many coefficients are of similar magnitude and thus ℓ_1 -minimization in any form is not very effective.

5.4. Resistor network

As our last example, consider the electrical resistor network shown in Figure 11. The network is comprised of $d = 2P$ resistances of uncertain ohmage and the network is driven by a voltage source providing a known potential V_0 . We are interested in determining how the voltage V shown in the figure depends on the $d = 2P$ resistances, which we take as random parameters uniformly distributed in the interval $\xi_k \in [1 - \varepsilon, 1 + \varepsilon]$, $k = 1, \dots, d$. This function is anisotropic. The effect of the resistors on the voltage will decay with distance (in terms of the number of preceding resistors) from the point V . In this example we set $d = 40$ ($P = 20$) and $d = 80$ ($P = 40$), take the maximum perturbation to be $\varepsilon = 0.1$ and set the reference potential $V_0 = 1$.

Figure 12 shows the error in the Legendre PCE for increasing design sizes M . In both cases the basis selection method produces a PCE that is significantly more accurate than the PCE produced by the non-adaptive strategy. The basis selection method provides comparable results to the approximately optimal oracle.

5.5. Gradient-enhanced ℓ_1 -minimization

Typical ℓ_1 -minimization, when used for PCE approximation, attempts to find solutions to

$$\Phi \alpha \approx \mathbf{f}$$

Figure 12: Convergence of the RMSE, with respect to increasing design size M , in the Legendre PCE approximation of the resistor network. The error bars (whiskers) represent the minimum and maximum error over the 20 trials for each design size M and the boxes represent the 1st, 2nd (median) and 3rd quartiles.

where denotes the Vandermonde matrix with entries $\Phi_{ij} = \phi_j(\xi_i)$, $i = 1, \dots, M$, $j = 1, \dots, N$. If gradients of the model f with respect to the random variables ξ are known, then one can enhance the accuracy of the PCE by finding a solution to

$$\Phi_{\partial} \alpha \approx \mathbf{f}_{\partial} \quad \text{where} \quad \Phi_{\partial} = \begin{bmatrix} \Phi \\ \frac{\partial \Phi}{\partial \xi_1} \\ \vdots \\ \frac{\partial \Phi}{\partial \xi_d} \end{bmatrix}, \quad \mathbf{f}_{\partial} = \begin{bmatrix} \mathbf{f} \\ \frac{\partial \mathbf{f}}{\partial \xi_1} \\ \vdots \\ \frac{\partial \mathbf{f}}{\partial \xi_d} \end{bmatrix}$$

and $(\frac{\partial \Phi}{\partial \xi_n})_{ij} = \frac{\partial \phi_j}{\partial \xi_n}(\xi_i)$ and $(\frac{\partial \mathbf{f}}{\partial \xi_n})_i = \frac{\partial f}{\partial \xi_n}(\xi_i)$, $i = 1, \dots, M$, $j = 1, \dots, N$, $n = 1, \dots, d$. To find a solution we again use basis pursuit denoising and solve

$$\alpha^* = \arg \min_{\alpha} \|\alpha\|_1 \quad \text{such that} \quad \|\Phi_{\partial} \alpha - \mathbf{f}_{\partial}\|_2 \leq \varepsilon$$

This gradient based formulation consists of $M(d+1)$ equations that match both function values and gradients, in comparison to (5) which consists of only M equations that match function values.

Figure 13 demonstrates the utility of using gradient data to build PCE approximations of the corner-peak function (11) with $d = 10$. Unlike the previous figures in this paper, the horizontal axis is no longer the number of model runs but rather the computational cost. We assume that running the model to only obtain function values costs one computational unit and running the model to obtain both function values and all gradients components

requires two units. For example, adjoint methods for differential equations can be used to obtain all gradients at a cost less than or equal to the cost of one forward model run.

Despite the extra computational cost required to obtain gradients, the use of gradients improves both the PCE resulting from both the non-adaptive and basis selection methods. Similar to the results presented in Section 5.1 the results shown here demonstrate that basis selection is more accurate than the non-adaptive strategy. Again the relative benefit is dependent on the rate of decay of the PCE coefficients.

Basis selection is able to make effective use of gradient information. For a design Ξ with M samples, the size of the gradient enhanced Vandermonde matrix is $M(d+1) \times N$. We see that for a given accuracy gradient-based PCE requires a factor of 4 fewer samples than the PCE based on the function values only. This is close to the optimal reduction factor of $d/2 = 5$ that can be obtained using gradients, assuming that each gradient component is as informative as a function value and the cost of computing function values with gradients is twice the cost of just computing function values.

5.6. Non-uniform model inputs

Throughout this paper, we have discussed basis selection when applied to Legendre polynomials and uniform variables. However, basis selection can also be applied to other variable/polynomial combinations. Let us once again consider the resistor network, but now let ξ be Gaussian variables with mean 1.0 and standard deviation 0.005.⁴ We now draw random samples from the aforementioned Gaussian distribution to form Ξ and run the model at each sample to obtain $f(\Xi)$. Figure 14 demonstrates that the advantages of basis selection are also present when we compute PCE approximations with non-uniform random variables.

Note that the error bars depicted in Figure 14 tend to increase in size with the number of samples M . This increase is likely caused by the poor numerical conditioning of the Vandermonde matrix in (5) which increases rapidly, with M , when using Hermite polynomials and normally distributed random variables. Development of sampling and pre-conditioning strategies for normal variables is an important area of future research, but is beyond the scope of this paper.

6. Conclusions

In this paper we present a basis selection method that can be used with ℓ_1 -minimization to adaptively determine the large coefficients of polynomial chaos expansions (PCE). The method attempts to identify structure in the coefficients of a PCE and only applies ℓ_1 -minimization to those terms believed to have large coefficients. The adaptive construction produces anisotropic basis sets that have more terms in important dimensions and limits the number of unimportant terms which increase mutual coherence and thus degrade the performance of ℓ_1 -minimization. The basis selection method produces, for a given computational budget, a more accurate PCE than would be obtained if the basis is fixed a priori. The important features and the accuracy of basis selection are demonstrated with a number of numerical examples. Specifically we show that in high dimensions,

⁴The standard deviation is made sufficiently small to make the chance of negative resistances practically zero.

(a)

(b)

(c)

Figure 13: Convergence of the RMSE, with respect to increasing computation units, in the gradient-enhanced Legendre PCE approximation of the corner-peak function (11) for the three coefficient regimes: (a) $\mathbf{c}^{(1)}$, (b) $\mathbf{c}^{(2)}$, (c) $\mathbf{c}^{(3)}$. The error bars (whiskers) represent the minimum and maximum error over the 20 trials for each design size M and the boxes represent the 1st, 2nd (median) and 3rd quartiles. We have assumed computing functional values at a design point is one unit and computing all gradient components is another unit. For example, the gradient enhanced basis selection approximation (dotted line) with 50 design points costs 100 units (horizontal axis value).

Figure 14: Convergence of the RMSE, with respect to increasing design size M , in the Hermite PCE approximation of the $d = 40$ resistor network. The error bars (whiskers) represent the minimum and maximum error over the 20 trials for each design size M and the boxes represent the 1st, 2nd (median) and 3rd quartiles.

for which high-order total-degree PCE bases are infeasible, basis selection allows accurate identification of high-order terms that cannot be captured by low-order total-degree expansions. We demonstrate that even for lower dimensional problems, for which high-order total-degree PCE bases are feasible, basis selection still produces more accurate results than basis sets that are fixed priori. Finally, we demonstrate that basis selection can effectively leverage function gradients and be applied to PCE of non-uniform random variables.

7. Acknowledgments

This work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy National Nuclear Security Administration under contract DE-AC04-94AL85000.

References

- [1] I. Babuska, R. Tempone, and G. Zouraris. Galerkin finite element approximations of stochastic elliptic partial differential equations. *SIAM Journal on Numerical Analysis*, 42(2):800–825, 2004.
- [2] I.M. Babuska, F. Nobile, and R. Tempone. A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 45(3):1005–1034, 2007.

- [3] R.G. Baraniuk, V. Cevher, M.F. Duarte, and C. Hegde. Model-based compressive sensing. *Information Theory, IEEE Transactions on*, 56(4):1982–2001, 2010.
- [4] J. Beck, R. Tempone, F. Nobile, and L. Tamellini. On the optimal polynomial approximation of stochastic pdes by galerkin and collocation methods. *Mathematical Models and Methods in Applied Sciences*, 22(09):1250023, 2012.
- [5] S. Becker, J. Bobin, and E. Candès. NESTA: A fast and accurate first-order method for sparse recovery. *SIAM Journal on Imaging Sciences*, 4(1):1–39, 2011.
- [6] Marcel Bieri and Christoph Schwab. Sparse high order {FEM} for elliptic spdes. *Computer Methods in Applied Mechanics and Engineering*, 198(1314):1149 – 1170, 2009. {HOFEM07} International Workshop on High-Order Finite Element Methods, 2007.
- [7] Géraud Blatman and Bruno Sudret. Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, 230(6):2345 – 2367, 2011.
- [8] Petros Boufounos, Marco F. Duarte, and Richard G. Baraniuk. Sparse signal reconstruction from noisy compressive measurements using cross validation. In *Statistical Signal Processing, 2007. SSP '07. IEEE/SP 14th Workshop on*, pages 299 –303, aug. 2007.
- [9] Leo Breiman and Philip Spector. Submodel selection and evaluation in regression. the x-random case. *International Statistical Review*, 60(3):pp. 291–319, 1992.
- [10] T.T. Cai and Lie Wang. Orthogonal matching pursuit for sparse signal recovery with noise. *Information Theory, IEEE Transactions on*, 57(7):4680–4688, July 2011.
- [11] E.J. Candes and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *Information Theory, IEEE Transactions on*, 52(12):5406–5425, Dec 2006.
- [12] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [13] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [14] S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM Rev.*, 43(1):129–159, January 2001.
- [15] A. Cohen, R. DeVore, and C. Schwab. Convergence rates of best n-term galerkin approximations for a class of elliptic spdes. *Foundations of Computational Mathematics*, 10(6):615–646, 2010.
- [16] Patrick R. Conrad and Youssef M. Marzouk. Adaptive smolyak pseudospectral approximations. *SIAM J. Scientific Computing*, 35(6), 2013.
- [17] P.G. Constantine, M.S. Eldred, and E.T. Phipps. Sparse pseudospectral approximation method. *Computer Methods in Applied Mechanics and Engineering*, 229232(0):1–12, 2012.
- [18] D.L. Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, April 2006.
- [19] D.L. Donoho, M. Elad, and V.N. Temlyakov. Stable recovery of sparse overcomplete representations in the presence of noise. *Information Theory, IEEE Transactions on*, 52(1):6–18, Jan 2006.
- [20] A. Doostan and H. Owhadi. A non-adapted sparse approximation of PDEs with stochastic inputs. *Journal of Computational Physics*, 230(8):3015 – 3034, 2011.
- [21] Marco F. Duarte, Michael B. Wakin, and Richard G. Baraniuk. Fast reconstruction of piecewise smooth signals from random projections. In *Online Proceedings of the Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, Rennes, France, 2005.
- [22] J. Foo and G.E. Karniadakis. Multi-element probabilistic collocation method in high dimensions. *J. Comput. Phys.*, 229(5):1536–1557, 2010.
- [23] B. Ganapathysubramanian and N. Zabaras. Sparse grid collocation schemes for stochastic natural convection problems. *Journal of Computational Physics*, 225(1):652–685, 2007.
- [24] A. Genz. A package for testing multiple integration subroutines. In P. Keas and G. Fairweather, editors, *Numerical Integration*, pages 337–340. D. Riedel, 1987.
- [25] R.G. Ghanem and P.D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer-Verlag New York, Inc., New York, NY, USA, 1991.
- [26] T. Hastie, R. Tibshirani, and J.H. Friedman. *The elements of statistical learning: data mining, inference, and prediction: with 200 full-color illustrations*. New York: Springer-Verlag, 2001.
- [27] J.D. Jakeman and S.G. Roberts. Local and dimension adaptive stochastic collocation for uncertainty quantification. In Jochen Garcke and Michael Griebel, editors, *Sparse Grids and Applications*, volume 88 of *Lecture Notes in Computational Science and Engineering*, pages 181–203. Springer Berlin Heidelberg, 2013.
- [28] Ron Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection.

- In *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI'95, pages 1137–1143, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc.
- [29] Stefan Kunis and Holger Rauhut. Random sampling of sparse trigonometric polynomials, ii. orthogonal matching pursuit versus basis pursuit. *Foundations of Computational Mathematics*, 8(6):737–763, 2008.
 - [30] C. La and M.N. Do. Tree-based orthogonal matching pursuit algorithm for signal reconstruction. In *Image Processing, 2006 IEEE International Conference on*, pages 1277–1280, 2006.
 - [31] X. Ma and N. Zabaras. An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations. *Journal of Computational Physics*, 228:3084–3113, 2009.
 - [32] L Mathelin and KA Gallivan. A compressed sensing approach for partial differential equations with random input data. *Commun. Comput. Phys.*, 12:919–954, 2012.
 - [33] L. Mathelin, M. Hussaini, and T. Zang. Stochastic approaches to uncertainty quantification in CFD simulations. *Numerical Algorithms*, 38(1-3):209–236, MAR 2005.
 - [34] Deanna Needell and Joel A. Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Commun. ACM*, 53(12):93–100, December 2010.
 - [35] F. Nobile, R. Tempone, and C.G. Webster. An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 46(5):2411–2442, 2008.
 - [36] Ji Peng, Jerrad Hampton, and Alireza Doostan. A weighted ℓ_1 -minimization approach for sparse polynomial chaos expansions. *Journal of Computational Physics*, 267(0):92 – 111, 2014.
 - [37] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.
 - [38] Holger Rauhut and Rachel Ward. Sparse legendre expansions via ℓ_1 -minimization. *Journal of Approximation Theory*, 164(5):517 – 533, 2012.
 - [39] Khachik Sargsyan, Cosmin Safta, Habib N Najm, Bert J Debusschere, Daniel Ricciuto, and Peter Thornton. Dimensionality reduction for complex models via Bayesian Compressive Sensing. *International Journal for Uncertainty Quantification*, 4(1):63–93, 2014.
 - [40] B. Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, 93(7):964–979, JUL 2008.
 - [41] M. Tatang, W. Pan, R. Prinn, and G. McRae. An efficient method for parametric uncertainty analysis of numerical geophysical model. *Journal of Geophysical Research*, 102(D18):21925–21932, 1997.
 - [42] R Tibshirani. Regression shrinkage and selection via the Lasso. *Journal Of The Royal Statistical Society Series B-Methodological*, 58(1):267–288, 1996.
 - [43] E. van den Berg and M.P. Friedlander. Probing the pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing*, 31(2):890–912, 2008.
 - [44] R. Ward. Compressed sensing with cross validation. *Information Theory, IEEE Transactions on*, 55(12):5773 –5782, dec. 2009.
 - [45] D. Xiu and J.S. Hesthaven. High-order collocation methods for differential equations with random inputs. *SIAM Journal on Scientific Computing*, 27(3):1118–1139, 2005.
 - [46] D. Xiu and G.E. Karniadakis. The Wiener-Askey Polynomial Chaos for stochastic differential equations. *SIAM J. Sci. Comput.*, 24(2):619–644, 2002.
 - [47] Z. Xu and T. Zhou. On sparse interpolation and the design of deterministic interpolation points. *SIAM Journal on Scientific Computing*, 2014. Accepted.
 - [48] L. Yan, L. Guo, and D. Xiu. Stochastic collocation algorithms using ℓ_1 -minimization. *International Journal for Uncertainty Quantification*, 2(3):279–293, 2012.
 - [49] X. Yang and G.E. Karniadakis. Reweighted minimization method for stochastic elliptic differential equations. *Journal of Computational Physics*, 248(0):87 – 108, 2013.