

A non-adapted sparse approximation of PDEs with stochastic inputs

Alireza Doostan^{a,*} Houman Owhadi^b

^a*Aerospace Engineering Sciences Department, University of Colorado, Boulder,
CO 80309, USA*

^b*Applied & Computational Mathematics Department, California Institute of
Technology, Pasadena, CA 91125, USA*

Abstract

We propose a method for the approximation of solutions of PDEs with stochastic coefficients based on the direct, i.e., non-adapted, sampling of solutions. This sampling can be done by using any legacy code for the deterministic problem as a black box. The method converges in probability (with probabilistic error bounds) as a consequence of sparsity and a concentration of measure phenomenon on the empirical correlation between samples. We show that the method is well suited for truly high-dimensional problems (with slow decay in the spectrum).

Key words: Polynomial chaos; Uncertainty quantification; Stochastic PDE; Compressive sampling; Sparse approximation

1 Introduction

Realistic analysis and design of complex engineering systems require not only a fine understanding and modeling of the underlying physics and their interactions, but also a significant recognition of intrinsic uncertainties and their influences on the quantities of interest. Uncertainty Quantification (UQ) is an emerging discipline that aims at addressing the latter issue; it aims at meaningful characterization of uncertainties in the physical models from the available measurements and efficient propagation of these uncertainties for a quantitative validation of model predictions.

* Corresponding author. Tel.: +(303)-492-7572; Fax: +(303)-492-4990.

Email addresses: doostan@colorado.edu (Alireza Doostan),
owhadi@caltech.edu (Houman Owhadi).

Despite recent growing interests in UQ of complex systems, it remains a grand challenge to efficiently propagate uncertainties through systems characterized by a large number of uncertain sources where the so-called curse-of-dimensionality is yet an unsolved problem. Additionally, development of *non-intrusive* uncertainty propagation techniques is of essence as the analysis of complex multi-disciplinary systems often requires the use of sophisticated coupled deterministic solvers in which one cannot readily intrude to set up the necessary propagation infrastructure.

Sampling methods such as the Monte Carlo simulation and its several variants had been utilized for a long time as the primary scheme for uncertainty propagation. However, it is well understood that these methods are generally inefficient for large-scale systems due to their slow rate of convergence. There has been an increasing recent interest in developing alternative numerical methods that are more efficient than the Monte Carlo techniques. Most notably, the stochastic Galerkin schemes using polynomial chaos (PC) bases [36, 27, 61, 2, 58] have been successfully applied to a variety of engineering problems and are extremely useful when the number of uncertain parameters is not large. In their original form, the stochastic Galerkin schemes are *intrusive*, as one has to modify the deterministic solvers for their implementation. *Stochastic collocation* schemes [53, 42, 60, 1, 47] belong to a different class of methods that rely upon (isotropic) sparse grid integration/interpolation in the stochastic space of the problem to reduce the curse-of-dimensionality associated with the conventional tensor-product integration/interpolation rules. As their construction is primarily based on the input parameter space, the computational cost of both stochastic Galerkin and collocation techniques increases rapidly for large number of independent input uncertainties.

More recently, efforts have been made to construct *solution-adaptive* uncertainty propagation techniques that exploit any structures in the solution to decrease the computational cost. Among them are the multi-scale model reduction of [32] and the sparse decomposition of [55, 8, 6, 7, 10] for the stochastic Galerkin technique, anisotropic and adaptive sparse grids of [46, 41] for the stochastic collocation scheme, and low-rank solution approximations of [48, 49, 33].

In the present study, we are interested in cases where the quantity of interest is *sparse* at the stochastic level, i.e., it can be accurately represented with only *few* terms when linearly expanded into a stochastic, e.g., polynomial chaos, basis. Interestingly, sparsity is salient in the analysis of high-dimensional problems where the number of energetic basis functions (those with large coefficients) is small relative to the cardinality of the full basis. For instance, it has been shown in [55, 8] that, under some mild conditions, solutions to linear elliptic stochastic PDEs with high-dimensional random coefficients admit sparse representations with respect to the PC basis. Consequently, an approach based

on a zero-dimensional algebraic stochastic problem has been proposed in [8] to detect the sparsity pattern, which then guides the stochastic Galerkin analysis of the original problem. Moreover, a “quasi”-best N -term approximation for a class of elliptic stochastic PDEs has been proposed in [7].

In this work, using *concentration of measure* inequalities and *compressive sampling* techniques, we derive a method for PC expansion of sparse solutions to stochastic PDEs. The proposed method is

- Non-intrusive: it is based on the direct random sampling of the PDE solutions. This sampling can be done by using any legacy code for the deterministic problem as a black box.
- Non-adapted: it does not tailor the sampling process to identify the important dimensions at the stochastic level
- Provably convergent: we obtain probabilistic bounds on the approximation error proving the stability and convergence of the method.
- Well-suited to problems with high-dimensional random inputs.

Compressive sampling is an emerging direction in signal processing that aims at recovering *sparse* signals accurately (or even exactly) from a small number of their random projections [19, 20, 16, 29, 14, 17, 15, 21, 13]. A sparse signal is simply a signal that has only few significant coefficients when linearly expanded into a basis, e.g., $\{\psi_\alpha\}$.

For sufficiently sparse signals, the number of samples needed for a successful recovery is typically less than what is required by the Shannon-Nyquist sampling principle. Generally speaking, a successful signal reconstruction by compressive sampling is conditioned upon:

- *Sufficient* sparsity of the signal; and
- *Incoherent* random projections of the signal.

A square-measurable stochastic function $u(\omega)$, defined on a suitable probability space $(\Omega, \mathcal{F}, \mathcal{P})$ can be expanded into a mean-squared convergent series of the chaos polynomial bases, i.e., $u(\omega) \approx \sum_\alpha c_\alpha \psi_\alpha(\omega)$, with some cardinality P . The stochastic function $u(\omega)$ is then sparse in PC basis $\{\psi_\alpha\}$, if only a small fraction of coefficients c_α are significant. In this case, under certain conditions, the sparse PC coefficients \mathbf{c} may be computed accurately and robustly using only $N \ll P$ random samples of $u(\omega)$ via compressive sampling. Given N random samples of $u(\omega)$, compressive sampling aims at finding the sparsest (or nearly sparsest) coefficients \mathbf{c} from an optimization problem of the form

$$(P_{s,\delta}) : \quad \min_{\mathbf{c}} \|\mathbf{W}\mathbf{c}\|_s \quad \text{subject to} \quad \|\Psi\mathbf{c} - \mathbf{u}\|_2 \leq \delta, \quad (1)$$

where $\|\mathbf{W}\mathbf{c}\|_s$, with $s = \{0, 1\}$ and some positive diagonal weight matrix \mathbf{W} , is a measure of the sparsity of \mathbf{c} and $\|\Psi\mathbf{c} - \mathbf{u}\|_2$ is a measure of the accuracy of

the truncated PC expansion in estimating the $u(\omega)$ samples. The N -vector \mathbf{u} contains the independent random samples of $u(\omega)$ and the rows of the $N \times P$ matrix Ψ consist of the corresponding samples of the PC basis $\{\psi_\alpha\}$.

Throughout the rest of this manuscript, we will elaborate on the formulation of the compressive sampling problem (1) and the required conditions under which it leads to an accurate and stable approximation of an arbitrary sparse stochastic function as well as sparse solutions to linear elliptic stochastic PDEs. Although we choose to study this particular class of stochastic PDEs, we stress that the proposed algorithms and theoretical developments are far more general and may be readily applied to recover sparse solution of other stochastic systems.

In Section 2, we describe the setup of the problem of interest. We then, in Section 3.1, briefly overview the spectral stochastic discretization of the random functions using PC basis. The main contribution of this work on sparse approximation of stochastic PDEs using compressive sampling is then introduced in Sections 3.2, 3.4, and 3.5. Sections 3.6 and 3.7 discuss some of the implementation details of the present technique. To demonstrate the accuracy and efficiency of the proposed procedures, in Section 4, we perform two numerical experiments on a 1- D (in space) linear elliptic stochastic PDE with high-dimensional random diffusion coefficients.

2 Problem setup

Let $(\Omega, \mathcal{F}, \mathcal{P})$ be a complete probability space where \mathcal{P} is a probability measure on the σ -field \mathcal{F} . We consider the following elliptic stochastic PDE defined on a bounded Lipschitz continuous domain $\mathcal{D} \subset \mathbb{R}^D$, $D = 1, 2$, or 3, with boundary $\partial\mathcal{D}$,

$$\begin{aligned} -\nabla \cdot (a(\mathbf{x}, \omega) \nabla u(\mathbf{x}, \omega)) &= f(\mathbf{x}) \quad \mathbf{x} \in \mathcal{D}, \\ u(\mathbf{x}, \omega) &= 0 \quad \mathbf{x} \in \partial\mathcal{D}, \end{aligned} \tag{2}$$

\mathcal{P} -*a.s.* $\omega \in \Omega$. The diffusion coefficient $a(\mathbf{x}, \omega)$ is a stochastic function defined on $(\Omega, \mathcal{F}, \mathcal{P})$ and is the source of uncertainty in (2). We assume that $a(\mathbf{x}, \omega)$ is specified by a truncated Karhunen-Loève-“like” expansion

$$a(\mathbf{x}, \omega) = \bar{a}(\mathbf{x}) + \sum_{i=1}^d \sqrt{\lambda_i} \phi_i(\mathbf{x}) y_i(\omega), \tag{3}$$

where (λ_i, ϕ_i) , $i = 1, \dots, d$, are the eigenpairs of the covariance function $C_{aa}(\mathbf{x}_1, \mathbf{x}_2) \in L_2(\mathcal{D} \times \mathcal{D})$ of $a(\mathbf{x}, \omega)$ and $\bar{a}(\mathbf{x})$ is the mean of $a(\mathbf{x}, \omega)$. We further assume that $a(\mathbf{x}, \omega)$ satisfies the following conditions:

A-I. For all $\mathbf{x} \in \mathcal{D}$, there exists constants a_{\min} and a_{\max} such that

$$0 < a_{\min} \leq a(\mathbf{x}, \omega) \leq a_{\max} < \infty \quad \mathcal{P} - a.s. \quad \omega \in \Omega \quad (4)$$

A-II. The covariance function $C_{aa}(\mathbf{x}_1, \mathbf{x}_2)$ is *piecewise analytic* on $\mathcal{D} \times \mathcal{D}$ [52, 8], implying that there exist real constants c_1 and c_2 such that for $i = 1, \dots, d$,

$$0 \leq \lambda_i \leq c_1 e^{-c_2 i^\kappa} \quad (5)$$

and

$$\forall \boldsymbol{\alpha} \in \mathbb{N}^d : \quad \sqrt{\lambda_i} \|\partial^{\boldsymbol{\alpha}} \phi_i\|_{L^\infty(\mathcal{D})} \leq c_1 e^{-c_2 i^\kappa}, \quad (6)$$

where $\kappa := 1/D$ and $\boldsymbol{\alpha} \in \mathbb{N}^d$ is a fixed multi-index. Notice that the decay rates in Eqs. (5) and (6) will be *algebraic* if $C_{aa}(\mathbf{x}_1, \mathbf{x}_2)$ has $C^s(\mathcal{D} \times \mathcal{D})$ regularity for some $s > 0$ [52].

A-III. The random variables $\{y_k(\omega)\}_{k=1}^d$ are independent and uniformly distributed on $\Gamma_k := [-1, 1]$, $k = 1, \dots, d$, with probability distribution function $\rho_k(y_k) = 1/2$ defined over Γ_k . The joint probability distribution function of the random vector $\mathbf{y} := (y_1, \dots, y_d)$ is then given by $\rho(\mathbf{y}) := \prod_{k=1}^d \rho_k(y_k)$.

Remark: The algorithm proposed in the paper requires the existence of a sparse solution. The only role of assumption A-II is to guarantee the existence of such a sparse solution. It is not necessary for the application and the validity of the proposed algorithm. In particular, if the coefficient a is only essentially bounded, the proposed algorithm will be accurate as long as a sparse approximation exists.

Given the finite-dimensional uncertainty representation in (3), the solution $u(\mathbf{x}, \omega)$ of (2) also admits a finite-dimensional representation, i.e.,

$$u(\mathbf{x}, \mathbf{y}) := u(\mathbf{x}, y_1(\omega), \dots, y_d(\omega)) : \mathcal{D} \times \Gamma \rightarrow \mathbb{R}, \quad (7)$$

where $\Gamma := \prod_{k=1}^d \Gamma_k$.

In what follows, we first briefly outline the Legendre spectral stochastic discretization of $u(\mathbf{x}, \mathbf{y})$ and consequently introduce our approach based on compressive sampling to obtain such a discretization.

3 Numerical approach

3.1 Spectral stochastic discretization

In the context of the spectral stochastic methods [36, 27, 61, 2], the solution $u(\mathbf{x}, \mathbf{y})$ of (2) is represented by an infinite series of the form

$$u(\mathbf{x}, \mathbf{y}) = \sum_{\boldsymbol{\alpha} \in \mathbb{N}_0^d} c_{\boldsymbol{\alpha}}(\mathbf{x}) \psi_{\boldsymbol{\alpha}}(\mathbf{y}), \quad (8)$$

where $\mathbb{N}_0^d := \{(\alpha_1, \dots, \alpha_d) : \alpha_j \in \mathbb{N} \cup \{0\}\}$ is the set of multi-indices of size d defined on non-negative integers. The basis functions $\{\psi_{\boldsymbol{\alpha}}(\mathbf{y})\}$ are multi-dimensional Legendre polynomials, referred to as the Legendre polynomial chaos (PC), and are orthogonal with respect to the joint probability measure $\rho(\mathbf{y})$ of the random vector \mathbf{y} . Each basis function $\psi_{\boldsymbol{\alpha}}(\mathbf{y})$ is a tensor product of univariate Legendre polynomials $\psi_{\alpha_i}(y_i)$ of degree $\alpha_i \in \mathbb{N}_0^1$, i.e.,

$$\psi_{\boldsymbol{\alpha}}(\mathbf{y}) = \psi_{\alpha_1}(y_1) \psi_{\alpha_2}(y_2) \cdots \psi_{\alpha_d}(y_d) \quad \boldsymbol{\alpha} \in \mathbb{N}_0^d. \quad (9)$$

We here assume that the univariate Legendre polynomials $\psi_{\alpha_i}(y_i)$ are also normalized such that

$$\int_{\Gamma_i} \psi_{\alpha_i}^2(y_i) \rho_i(y_i) dy_i = 1, \quad i = 1, \dots, d. \quad (10)$$

The exact generalized Fourier coefficients $c_{\boldsymbol{\alpha}}(\mathbf{x})$ in (8), referred to as the PC coefficients, are computed by the projection of $u(\mathbf{x}, \mathbf{y})$ onto each basis function $\psi_{\boldsymbol{\alpha}}(\mathbf{y})$,

$$c_{\boldsymbol{\alpha}}(\mathbf{x}) = \mathbb{E}[u(\mathbf{x}, \mathbf{y}) \psi_{\boldsymbol{\alpha}}(\mathbf{y})] = \int_{\Gamma} u(\mathbf{x}, \mathbf{y}) \psi_{\boldsymbol{\alpha}}(\mathbf{y}) \rho(\mathbf{y}) d\mathbf{y}. \quad (11)$$

Here, \mathbb{E} denotes the expectation operator. In practice, the expansion (8) is finite; that is, only a finite number of spectral modes is needed to approximate $u(\mathbf{x}, \mathbf{y})$ within a given target accuracy. Traditionally, a finite-order truncation of the basis $\{\psi_{\boldsymbol{\alpha}}(\mathbf{y})\}$ is adopted for the approximation, i.e.,

$$u_p(\mathbf{x}, \mathbf{y}) := \sum_{\boldsymbol{\alpha} \in \Lambda_{p,d}} c_{\boldsymbol{\alpha}}(\mathbf{x}) \psi_{\boldsymbol{\alpha}}(\mathbf{y}), \quad (12)$$

where the set of multi-indices $\Lambda_{p,d}$ is

$$\Lambda_{p,d} := \left\{ \boldsymbol{\alpha} \in \mathbb{N}_0^d : \|\boldsymbol{\alpha}\|_1 \leq p, \|\boldsymbol{\alpha}\|_0 \leq d \right\} \quad (13)$$

and has the cardinality

$$P := |\Lambda_{p,d}| = \frac{(p+d)!}{p!d!}. \quad (14)$$

Here, $\|\boldsymbol{\alpha}\|_1 = \sum_{i=1}^d \alpha_i$ and $\|\boldsymbol{\alpha}\|_0 = \#\{i : \alpha_i > 0\}$ are the total order (degree) and dimensionality of the basis function $\psi_{\boldsymbol{\alpha}}(\mathbf{y})$, respectively. The approximation is then refined by increasing p to achieve a given target accuracy. Under assumptions *A-I*, *A-II*, and *A-III* stated in Section 2, the solution $u(\mathbf{x}, \mathbf{y})$ is analytic with respect to the random variables $\{y_i\}_{i=1}^d$ (see [1]), and as p increases,

the approximation (12) converges exponentially fast in the mean-squares sense [2, 1, 8].

Definition (Sparsity) *The solution $u(\mathbf{x}, \mathbf{y})$ is said to be (nearly) sparse if only a small fraction of coefficients $c_\alpha(\mathbf{x})$ in (12) are dominant and contribute to the solution statistics.*

As will be described in Section 3.2, a sparse solution $u(\mathbf{x}, \mathbf{y})$ may be accurately recovered using $N \ll P$ random samples $\{u(\mathbf{x}, \mathbf{y}_i)\}_{i=1}^N$ using compressive sampling techniques. This has to be compared, for instance, with the least-squares regression-type techniques, [38], that normally require $N \gg P$ samples for an accurate recovery.

3.2 Sparse recovery using compressive sampling

Compressive sampling is an emerging theory in the field of signal and image processing [19, 20, 16, 29, 14, 17, 15, 21, 13]. It hinges around the idea that a set of incomplete random observations of a sparse signal can be used to accurately, or even exactly, recover the signal (provided that the basis in which the signal is sparse is known). In particular, the number of such observations may be much smaller than the cardinality of the signal. In the context of problem (2), compressive sampling may be interpreted as follows. The solution $u(\mathbf{x}, \mathbf{y})$ that is sparse, in the sense of Lemma 3.2 defined in Section 3.3, can be accurately recovered using $N \ll P$ random samples $\{u(\mathbf{x}, \mathbf{y}_i)\}_{i=1}^N$, where P is the cardinality of the Legendre PC basis $\{\psi_\alpha(\mathbf{y})\}$. We next elaborate on the above statement and address how such a sparse reconstruction is achieved and under what conditions it is successful.

Let $\{u(\mathbf{y}_i)\}_{i=1}^N$ be *i.i.d.* random samples of $u(\mathbf{x}, \mathbf{y})$ for a fixed point \mathbf{x} in \mathcal{D} . For the time being, let us assume that the p th-order PC basis $\{\psi_\alpha(\mathbf{y})\}$ is a complete basis to expand $u(\mathbf{y})$; we will relax this assumption as we proceed. Given pairs of $\{\mathbf{y}_i\}_{i=1}^N$ and $\{u(\mathbf{y}_i)\}_{i=1}^N$, we write

$$u(\mathbf{y}_i) = \sum_{\alpha \in \Lambda_{p,d}} c_\alpha \psi_\alpha(\mathbf{y}_i), \quad i = 1, \dots, N, \quad (15)$$

or equivalently,

$$\mathbf{\Psi} \mathbf{c} = \mathbf{u}. \quad (16)$$

Here $\mathbf{c} \in \mathbb{R}^P$ is the vector of PC coefficients c_α to be determined, $\mathbf{u} \in \mathbb{R}^N$ is the vector of samples of $u(\mathbf{y})$, and each column of the *measurement* matrix $\mathbf{\Psi} \in \mathbb{R}^{N \times P}$ contains samples of the j th element of the PC basis, i.e.,

$$\Psi[i, j] = \psi_{\alpha_j}(\mathbf{y}_i), \quad i = 1, \dots, N, \quad j = 1, \dots, P. \quad (17)$$

We are interested in the case that the number N of solution samples is much smaller than the unknown PC coefficients P , i.e., $N \ll P$. Without any additional constraints on \mathbf{c} , the underdetermined linear system (16) is ill-posed and, in general, has infinitely many solutions. When \mathbf{c} is sparse; that is, only a small fraction of the coefficients c_α are significant, the problem (16) may be regularized to ensure a well-posed solution. Such a regularization may be achieved by seeking a solution \mathbf{c} with the minimum number of non-zeros. This can be formulated in the optimization problem

$$(P_0) : \quad \min_{\mathbf{c}} \|\mathbf{c}\|_0 \quad \text{subject to} \quad \Psi \mathbf{c} = \mathbf{u}, \quad (18)$$

where the semi-norm $\|\mathbf{c}\|_0 := \#\{\alpha : c_\alpha \neq 0\}$ is the number of non-zero components of \mathbf{c} . In general, the global minimum solution of (P_0) is not unique and is NP-hard to compute: the cost of a global search is exponential in P . Further developments in compressive sampling resulted in a convex relaxation of problem (P_0) by minimization of the ℓ_1 -norm of the solution \mathbf{c} instead, i.e.,

$$(P_1) : \quad \min_{\mathbf{c}} \|\mathbf{W}\mathbf{c}\|_1 \quad \text{subject to} \quad \Psi \mathbf{c} = \mathbf{u}, \quad (19)$$

where \mathbf{W} is a diagonal matrix whose $[j, j]$ entry is the ℓ_2 -norm of the j th column of Ψ and $\|\cdot\|_1$ denotes the ℓ_1 -norm. Notice that the ℓ_1 -norm is the closest convex function to the ℓ_0 -norm that compels the small coefficients c_α to be zero, thus promoting the sparsity in the solution. The purpose of weighting the ℓ_1 cost function with \mathbf{W} is to prevent the optimization from biasing toward the non-zero entries in \mathbf{c} whose corresponding columns in Ψ have large norms. The problem (P_1) is entitled Basis Pursuit (BP) [19] and its solution can be obtained by linear programming. Since the ℓ_1 -norm functional $\|\mathbf{c}\|_1$ is convex, the optimization problem (P_1) admits a unique solution that coincides with the unique solution to problem (P_0) for sufficiently sparse \mathbf{c} with some constraints on the measurement matrix Ψ ; e.g., see [13].

In general, the p th-order PC basis is not complete for the exact representation of $u(\mathbf{y})$; therefore, we have to account for the truncation error. This can be accommodated in (P_0) and (P_1) by allowing a non-zero residual in the constraint $\Psi \mathbf{c} = \mathbf{u}$. Therefore, as in Sections 3.2.1 and 3.2.3 of [13], the proposed algorithms in this paper are error-tolerant versions of (P_0) and (P_1) , with error tolerance δ , i.e.,

$$(P_{0,\delta}) : \quad \min_{\mathbf{c}} \|\mathbf{c}\|_0 \quad \text{subject to} \quad \|\Psi \mathbf{c} - \mathbf{u}\|_2 \leq \delta \quad (20)$$

and

$$(P_{1,\delta}) : \quad \min_{\mathbf{c}} \|\mathbf{W}\mathbf{c}\|_1 \quad \text{subject to} \quad \|\Psi \mathbf{c} - \mathbf{u}\|_2 \leq \delta, \quad (21)$$

respectively. The latter problem is named Basis Pursuit Denoising (BPDN) in [19] and may be solved using techniques from quadratic programming. We leave the discussion on the available algorithms for solving problems $(P_{1,\delta})$ and

$(P_{0,\delta})$ to Section 3.7. Instead, we henceforth delineate on sufficient conditions under which the BPDN problem $(P_{1,\delta})$ leads to a successful Legendre PC expansion of a general essentially bounded sparse stochastic function $u(\mathbf{y})$ and, subsequently, the sparse solution $u(\mathbf{x}, \mathbf{y})$ to the problem (2). Our results are extensions of those in [31, 13], adapted to the case where the measurement matrix Ψ consists of random evaluations of the Legendre PC basis $\{\psi_\alpha\}$. With slight differences that will be remarked accordingly, similar results hold for the case of the $(P_{0,\delta})$ problem.

Theorem 3.1 (General stability of $(P_{1,\delta})$) *Let $u(\mathbf{y})$ be an essentially bounded function of i.i.d. random variables $\mathbf{y} := (y_1, \dots, y_d)$ uniformly distributed on $\Gamma := [-1, 1]^d$. Define*

$$S_{\max} := \frac{N}{64P^{4c_{p,d}}(\ln P)}, \quad (22)$$

with

$$c_{p,d} := \frac{\ln 3}{2} \frac{p}{\ln\left(\frac{(p+d)!}{p!d!}\right)}. \quad (23)$$

Let $u_p^{1,\delta}(\mathbf{y}) := \sum_{\alpha \in \Lambda_{p,d}} c_\alpha^{1,\delta} \psi_\alpha(\mathbf{y})$ be the Legendre PC approximation of $u(\mathbf{y})$ with coefficients $\mathbf{c}^{1,\delta}$ computed from the ℓ_1 -minimization problem $(P_{1,\delta})$ with $\delta \geq 0$. If there exists a Legendre PC expansion $u_p^0(\mathbf{y}) := \sum_{\alpha \in \Lambda_{p,d}^\epsilon} c_\alpha^0 \psi_\alpha(\mathbf{y})$, for some index set $\Lambda_{p,d}^\epsilon \subseteq \Lambda_{p,d}$ such that $\|u - u_p^0\|_{L^\infty(\Gamma)} \leq \epsilon$ and

$$S < S_{\max}, \quad (24)$$

with $S := |\Lambda_{p,d}^\epsilon|$, then with probability

$$Prob_1 \geq 1 - 4P^{2-2S_{\max}} - P^{-8S_{\max}} - P^{-8S_{\max}P^{4c_{p,d}}}, \quad (25)$$

(on the N samples $\{u(\mathbf{y}_i)\}_{i=1}^N$) and for some constants c_1 and c_2 , the solution $u_p^{1,\delta}$ must obey

$$\|u - u_p^{1,\delta}\|_{L^2(\Gamma)} \leq c_1\epsilon + c_2 \frac{\delta}{\sqrt{N}}. \quad (26)$$

In simple words, Theorem 3.1 states that if an essentially bounded stochastic function $u(\mathbf{y})$ admits a sparse Legendre PC expansion, then the ℓ_1 -minimization problem $(P_{1,\delta})$ can accurately recover it from a sufficiently large number of random solution samples. The recovery is stable under the truncation error $\|\Psi\mathbf{c} - \mathbf{u}\|_2$ and is within a distance of the exact solution that is proportional to the error tolerance δ . It is worth highlighting that no prior knowledge of the sparsity pattern of the PC coefficients \mathbf{c} is needed for an accurate recovery.

Remark: Based on the conditions (22) and (24), the number N of random samples has to grow like $P^{4c_{p,d}} \ln P$ and also proportional to the number of dominant coefficients $S = |\Lambda_{p,d}^\epsilon|$. Given any order p of the PC expansion, for sufficiently high-dimensional problems, the constant $c_{p,d} < 1/4$ (see Lemma

3.5 and Fig. 1), thus justifying $N \ll P$. In fact, the conditions (22) and (24) are too pessimistic; in practice, the number of random samples required for an accurate recovery is much smaller than the theoretical value in (22). We will elaborate on this statement in Section 3.5.

Remark: Although the BPDN reconstruction is achieved by minimizing the ℓ_1 -norm of the solution, based on (26), the approximation also converges to the exact solution in the mean-squares sense.

Remark: A similar theorem holds for the case of the sparse approximation using the ℓ_0 -minimization problem $(P_{0,\delta})$ in (20). In this case, the condition (22) has to be replaced with $S_{\max} := \frac{N}{16P^{4c_{p,d}(\ln P)}}$ which is, in theory, milder than that of the $(P_{1,\delta})$ problem. The error estimate (26) also holds with a larger probability, but with different constants c_1 and c_2 .

3.3 Sparsity of the solution $u(\mathbf{x}, \mathbf{y})$

Notice that the accurate recovery of $u(\mathbf{y})$ is conditioned upon the existence of a sparse PC expansion u_p^0 (see Theorem 3.1). In fact, this assumption may not hold for an arbitrary stochastic function $u(\mathbf{y})$, as all the elements of the basis set $\{\psi_\alpha(\mathbf{y})\}$ may be important. In this case, our sparse approximation still converges to the actual solution but, perhaps, not using as few as $N \ll P$ random solution samples.

We will now summarize the results of [55, 8] on the sparsity of the Legendre PC expansion of the solution $u(\mathbf{x}, \mathbf{y})$ to the problem (2). Alternative to the p th-order truncated PC expansion of (12), one may ideally seek a proper index set $\Lambda_{p,d}^\epsilon \subseteq \Lambda_{p,d}$, with sufficiently large p , such that for a given accuracy ϵ

$$\Lambda_{p,d}^\epsilon := \arg \min \left\{ |\tilde{\Lambda}_{p,d}| : \tilde{\Lambda}_{p,d} \subseteq \Lambda_{p,d}, \|u - \tilde{u}_p\|_{H_0^1(\mathcal{D}, L^\infty(\Gamma))} \leq \epsilon \right\}, \quad (27)$$

in which $\tilde{u}_p(\mathbf{x}, \mathbf{y}) := \sum_{\alpha \in \tilde{\Lambda}_{p,d}} c_\alpha(\mathbf{x}) \psi_\alpha(\mathbf{y})$. This will lead to the so-called *sparse approximation* of $u(\mathbf{x}, \mathbf{y})$ if

$$|\Lambda_{p,d}^\epsilon| \ll |\Lambda_{p,d}| = P, \quad (28)$$

where $\Lambda_{p,d}$ is defined in (13). Such a reduction in the number of basis functions in (28) is possible as, given the accuracy ϵ , the *effective* dimensionality ν of $u(\mathbf{x}, \mathbf{y})$ in Γ is potentially smaller than the apparent dimensionality d . More precisely, under assumptions *A-I*, *A-II*, and *A-III* stated in Section 2, the analyses of [55, 8] imply that the discretization of $u(\mathbf{x}, \mathbf{y})$ using a sparse index set $\Lambda_{p,\nu}$,

$$\Lambda_{p,\nu} := \left\{ \boldsymbol{\alpha} \in \mathbb{N}_0^d : \|\boldsymbol{\alpha}\|_1 \leq p, \|\boldsymbol{\alpha}\|_0 \leq \nu \leq d \right\} \quad (29)$$

preserves the exponential decay of the approximation error in the $H_0^1(\mathcal{D}, L^\infty(\Gamma))$ sense. For the sake of completeness, we cite this from [8] in the following lemma.

Lemma 3.2 (Proposition 3.10 of [8]) *Given assumptions A-I, A-II, and A-III in Section 2, there exist constants $c_1, c_2, c_3, c_4 > 0$, depending only on $a(\mathbf{x}, \omega)$ and $f(\mathbf{x})$ but independent of d, p, ν , such that*

$$\|u - u_{p,\nu}\|_{H_0^1(\mathcal{D}, L^\infty(\Gamma))} \leq c_1 \left(e^{-c_2\nu^{1+\kappa}} + e^{c_3\nu(\ln d + \ln p) - c_4 p} \right), \quad (30)$$

for any $d, p, \nu \in \mathbb{N}$ with $\nu \leq d$ and $\kappa = 1/D$.

In particular, for $d \geq c_d |\ln \epsilon|^{1/\kappa}$, choosing

$$p_\epsilon = \lceil c_p d^\kappa \rceil \leq p \quad \text{and} \quad \nu_\epsilon = \lceil c_\nu d^{\kappa/(\kappa+1)} \rceil \leq d, \quad (31)$$

leads to

$$\|u - u_{p_\epsilon, \nu_\epsilon}\|_{H_0^1(\mathcal{D}, L^\infty(\Gamma))} \leq \epsilon, \quad (32)$$

where $u_{p_\epsilon, \nu_\epsilon}$ is now defined on a sparse index set

$$\Lambda_{p_\epsilon, \nu_\epsilon} := \left\{ \boldsymbol{\alpha} \in \mathbb{N}_0^d : \|\boldsymbol{\alpha}\|_1 \leq p_\epsilon, \|\boldsymbol{\alpha}\|_0 \leq \nu_\epsilon \right\} \quad (33)$$

with cardinality

$$|\Lambda_{p_\epsilon, \nu_\epsilon}| \lesssim \epsilon^{-1/\rho}, \quad (34)$$

for some arbitrary large $\rho > 0$ and constants c_d, c_p , and c_ν independent of d, p_ϵ , and ν_ϵ [8].

In practice, the sparse set $\Lambda_{p,d}^\epsilon$ in (27) (or equivalently $\Lambda_{p_\epsilon, \nu_\epsilon}$ in (33)) is not known *a priori*. In [8], an approach based on an algebraic purely-stochastic problem is proposed to adaptively identify $\Lambda_{p,d}^\epsilon$. Having done this, the coefficients of the the spectral modes are computed via the (intrusive) stochastic Galerkin scheme [36, 61]. Alternatively, in this work, we apply our sparse approximation using $(P_{1,\delta})$ and $(P_{0,\delta})$ to compute $u(\mathbf{x}, \mathbf{y})$. The implementation of $(P_{1,\delta})$ and $(P_{0,\delta})$ is non-intrusive; only random samples of the solution are needed. Moreover, we do not adapt the sampling process to identify the important dimensions at the stochastic level; therefore, our constructions are non-adapted.

Throughout the rest of the present paper, we focus our attention on the case of the stochastic PDE (2) whose solution is provably sparse. The statement of Theorem 3.1 can be specialized to the approximation of the sparse solution to the stochastic PDE (2) using $(P_{1,\delta})$ (or $(P_{0,\delta})$) as follows.

3.4 Stability of $(P_{1,\delta})$ for stochastic PDE (2)

Combining Lemma 3.2 with Theorem 3.1 leads to the following theorem.

Theorem 3.3 (Stability of $(P_{1,\delta})$ for stochastic PDE (2)) *Let $u_p^{1,\delta}(\mathbf{x}, \mathbf{y}) := \sum_{\alpha \in \Lambda_{p,d}} c_{\alpha}^{1,\delta}(\mathbf{x}) \psi_{\alpha}(\mathbf{y})$ be the p th-order Legendre PC approximation of $u(\mathbf{x}, \mathbf{y})$ in (2) where the coefficients $c_{\alpha}^{1,\delta}(\mathbf{x})$ are obtained from the ℓ_1 -minimization problem $(P_{1,\delta})$ with N independent samples of $u(\mathbf{x}, \mathbf{y})$ and arbitrary $\delta \geq 0$. Write $\kappa := 1/D$. Let $c_{p,d}$ be defined by (23) and let S_{\max} be defined by (22). Let $\rho > 0$ be arbitrary.*

There exists constants c_1, c_2, c_3, c_4, c_5 independent from p, d, κ, N such that if $\lceil c_2 d^{\kappa} \rceil \leq p$ and $\lceil c_3 d^{\kappa/(\kappa+1)} \rceil \leq d$,

then with probability at least

$$\text{Prob}_1 \geq 1 - 4P^{2-2S_{\max}} - P^{-8S_{\max}} - P^{-8P^{4c_{p,d}}S_{\max}}, \quad (35)$$

the solution $u_p^{1,\delta}$ must obey

$$\|u - u_p^{1,\delta}\|_{L^2(\mathcal{D}, L^2(\Gamma))} \leq c_4 \epsilon + c_5 \frac{\delta}{\sqrt{N}}. \quad (36)$$

with

$$\epsilon := \max \left(\frac{1}{S_{\max}^{\rho}}, \exp \left(- \left(\frac{d}{c_1} \right)^{\kappa} \right) \right) \quad (37)$$

3.5 Proofs and further ingredients of successful sparse approximations via $(P_{1,\delta})$ and $(P_{0,\delta})$

The ability of problems $(P_{1,\delta})$ and $(P_{0,\delta})$ in accurately approximating the sparse PC coefficients \mathbf{c} in (12), hence the solution $u(\mathbf{x}, \mathbf{y})$, depends on two main factors: *i*) the sparsity of the PC coefficients \mathbf{c} and *ii*) the *mutual coherence* of the measurement matrix Ψ . In fact, the number N of random solution samples required for a successful sparse approximation is dictated by these two factors. While sparsity is a characteristic of the solution of interest $u(\mathbf{x}, \mathbf{y})$, the mutual coherence of the measurement matrix Ψ is universal as it only depends on the choice of PC basis $\{\psi_{\alpha}(\mathbf{y})\}$ and the sampling process from which Ψ is assembled.

In Section 3.3, based on the analysis of [8], we rationalized the sparsity of $u(\mathbf{x}, \mathbf{y})$ with respect to the Legendre PC basis. We now give the definition of the mutual coherence of Ψ and discuss its role in our sparse approximation using $(P_{1,\delta})$ and $(P_{0,\delta})$.

3.5.1 Mutual coherence of Ψ

Definition (Mutual Coherence [31]) *The mutual coherence $\mu(\Psi)$ of a matrix $\Psi \in \mathbb{R}^{N \times P}$ is the maximum of absolute normalized inner-products of its columns. Let ψ_j and ψ_k be two distinct columns of Ψ . Then,*

$$\mu(\Psi) := \max_{1 \leq j, k \leq P, j \neq k} \frac{|\psi_j^T \psi_k|}{\|\psi_j\|_2 \|\psi_k\|_2}. \quad (38)$$

In plain words, the mutual coherence is a measure of how close to orthogonal a matrix is. Clearly, for any general matrix Ψ ,

$$0 \leq \mu(\Psi) \leq 1, \quad (39)$$

where the lower bound is achieved, for instance, by unitary matrices. However, for the case of $N < P$, the mutual coherence $\mu(\Psi)$ is strictly positive. It is well understood that measurement matrices with smaller mutual coherence have a better ability to recover a sparse solution using compressive sampling techniques, e.g., see Lemma 3.6. Therefore, we shall proceed to examine the mutual coherence of the random measurement matrix Ψ in (16). We first observe that, by the orthogonality of the Legendre PC basis, the mutual coherence $\mu(\Psi)$ converges to zero almost surely for asymptotically large random sample sizes N . However, it is essential for our purpose to *i*) investigate if a desirably small $\mu(\Psi)$ can be achieved by a sample size $N \ll P$ and *ii*) quantify how large $\mu(\Psi)$ can get for a *finite* N . These are addressed in the following theorem.

Theorem 3.4 (Bound on the mutual coherence $\mu(\Psi)$) *Let $\Psi \in \mathbb{R}^{N \times P}$, as defined in (16), be the measurement matrix corresponding to N independent random samples of the Legendre polynomial chaos basis of order p in d i.i.d. uniform random variables \mathbf{y} . There exists a positive constant $c_{p,d} := \frac{\ln 3}{2} \frac{p}{\ln\left(\frac{(p+d)!}{p!d!}\right)}$ depending on p and d , such that if*

$$0 \leq r = 2\sqrt{\zeta P^{4c_{p,d}}(\ln P)/N} \leq 1/2, \quad (40)$$

for some $\zeta > 1$, then

$$\text{Prob} \left[\mu(\Psi) \geq \frac{r}{1-r} \right] \leq 4P^{2-2\zeta}. \quad (41)$$

Figure 1 illustrates the decay of $c_{p,d}$, for several values of p , as a function of d . Based on Theorem 3.4, for cases where the number d of random variables \mathbf{y} is large enough such that $c_{p,d} < 1/4$, it is sufficient to have $N \sim \mathcal{O}(16P^{4c_{p,d}} \ln P) \ll P$ to keep $\mu(\Psi)$ bounded from above with a large probability. Notice that such a requirement on $c_{p,d}$ is particularly suited to high-dimensional problems.

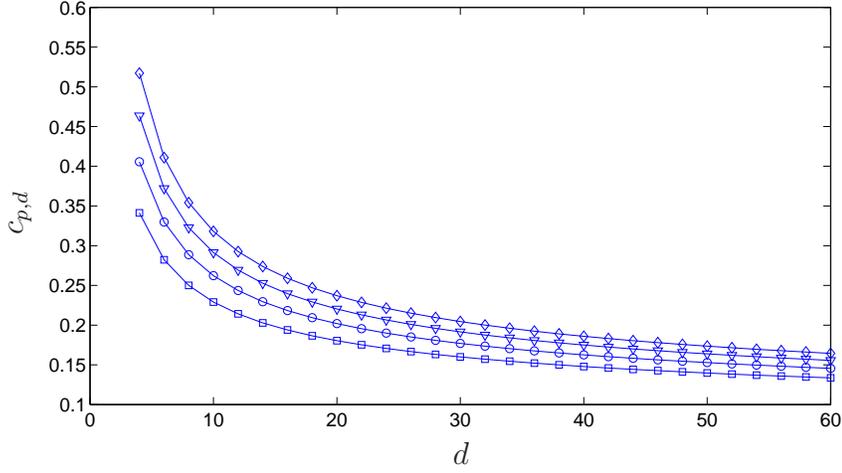


Fig. 1. Decay of $c_{p,d}$ as a function of d . $p = 1$ (□); $p = 2$ (○); $p = 3$ (▽); $p = 4$ (◇).

Remark: We observe that, given the choice of r in (40), the upper bound on $\mu(\Psi)$ in (41) decays like $1/\sqrt{N}$ for asymptotically large N , which is consistent with the Central Limit Theorem.

In order to prove Theorem 3.4, we first need to compute the maximum of the Legendre PC basis functions $\psi_{\alpha}(\mathbf{y})$. This is given in the following lemma.

Lemma 3.5 (Maximum of $\psi_{\alpha}(\mathbf{y})$) *Let $\{\psi_{\alpha}(\mathbf{y})\}$ be the Legendre polynomial chaos basis of total order p in d i.i.d uniform random variables \mathbf{y} (as defined in (13)) and with cardinality P . Then,*

$$\|\psi_{\alpha}\|_{L^{\infty}(\Gamma)} \leq P^{c_{p,d}}, \quad (42)$$

with a constant

$$c_{p,d} := \frac{\ln 3}{2} \frac{p}{\ln \left(\frac{(p+d)!}{p!d!} \right)}. \quad (43)$$

Proof. Given the equality

$$\|\psi_{\alpha_i}\|_{L^{\infty}(\Gamma_i)} = \sqrt{2\alpha_i + 1},$$

we have

$$\|\psi_{\alpha}\|_{L^{\infty}(\Gamma)} = \prod_{i=1}^d \sqrt{2\alpha_i + 1}.$$

Using the constraint $\alpha_i \in \mathbb{N}_0^1$ and $\sum_{i=1}^d \alpha_i \leq p$, the right-hand-side is maximized when p of α_i s are equal to one and $d - p$ equal to zero (we assume $d \geq p$). We deduce that

$$\|\psi_{\alpha}\|_{L^{\infty}(\Gamma)} \leq 3^{\frac{p}{2}} = P^{\frac{\ln 3}{2} \frac{p}{\ln P}} = P^{c_{p,d}}.$$

We are now ready to prove Theorem 3.4.

Proof of Theorem 3.4. The mutual coherence $\mu(\Psi)$ is

$$\mu(\Psi) = \max_{1 \leq j, k \leq P, j \neq k} \frac{\left| \frac{1}{N} \sum_{i=1}^N \psi_{\alpha_j}(\mathbf{y}_i) \psi_{\alpha_k}(\mathbf{y}_i) \right|}{\left(\frac{1}{N} \sum_{i=1}^N \psi_{\alpha_j}^2(\mathbf{y}_i) \right)^{1/2} \left(\frac{1}{N} \sum_{i=1}^N \psi_{\alpha_k}^2(\mathbf{y}_i) \right)^{1/2}}. \quad (44)$$

Given the independence of samples $\{\mathbf{y}_i\}_{i=1}^N$ and using the McDiarmid's inequality, we obtain

$$\text{Prob} \left[\left| \frac{1}{N} \sum_{i=1}^N \psi_{\alpha_j}(\mathbf{y}_i) \psi_{\alpha_k}(\mathbf{y}_i) \right| \geq r \right] \leq 2 \exp \left(\frac{-2Nr^2}{(2\|\psi_{\alpha_j}(\mathbf{y}_i) \psi_{\alpha_k}(\mathbf{y}_i)\|_{L^\infty(\Gamma_i)})^2} \right). \quad (45)$$

Using Lemma 3.5, we may probabilistically bound the numerator in (44) as

$$\text{Prob} \left[\left| \frac{1}{N} \sum_{i=1}^N \psi_{\alpha_j}(\mathbf{y}_i) \psi_{\alpha_k}(\mathbf{y}_i) \right| \geq r \right] \leq 2 \exp \left(\frac{-Nr^2}{2P^{4c_{p,d}}} \right),$$

in which we exploit the orthonormality of $\psi_{\alpha_j}(\mathbf{y})$ and $\psi_{\alpha_k}(\mathbf{y})$, i.e., $\mathbb{E}[\psi_{\alpha_j} \psi_{\alpha_k}] = \delta_{jk}$. Similarly, for $j = 1, \dots, P$, we have

$$\text{Prob} \left[\frac{1}{N} \sum_{i=1}^N \psi_{\alpha_j}^2(\mathbf{y}_i) \leq 1 - r \right] \leq \exp \left(\frac{-Nr^2}{P^{4c_{p,d}}} \right) \leq \exp \left(\frac{-Nr^2}{2P^{4c_{p,d}}} \right).$$

Therefore,

$$\text{Prob} \left[\frac{\left| \frac{1}{N} \sum_{i=1}^N \psi_{\alpha_j}(\mathbf{y}_i) \psi_{\alpha_k}(\mathbf{y}_i) \right|}{\left(\frac{1}{N} \sum_{i=1}^N \psi_{\alpha_j}^2(\mathbf{y}_i) \right)^{1/2} \left(\frac{1}{N} \sum_{i=1}^N \psi_{\alpha_k}^2(\mathbf{y}_i) \right)^{1/2}} \geq \frac{r}{1-r} \right] \leq 4 \exp \left(\frac{-Nr^2}{2P^{4c_{p,d}}} \right) \quad (46)$$

and

$$\text{Prob} \left[\mu(\Psi) \geq \frac{r}{1-r} \right] \leq 4P^2 \exp \left(\frac{-Nr^2}{2P^{4c_{p,d}}} \right). \quad (47)$$

Taking

$$r = 2\sqrt{\zeta P^{4c_{p,d}}(\ln P)/N}, \quad (48)$$

for some $\zeta > 1$, we arrive at the statement of the Theorem 3.4. \square

To summarize, we observe that with large probability, the mutual coherence $\mu(\Psi)$ of the measurement matrix Ψ in (16) can be arbitrarily bounded from above by increasing the number N of independent random solution samples. Moreover, given the discussions of Section 3.3, we know that the solution to problem (2) is sparse in the Legendre PC basis. These are the two key factors affecting the stability and accuracy of our sparse approximation.

Following [31, 13], we next state a *sufficient* condition on the sparsity of $u(\mathbf{x}, \mathbf{y})$ (or, equivalently, the mutual coherence of Ψ) such that the problem $(P_{1,\delta})$ leads to a stable and accurate sparse solution. By stability, we simply mean that

the PC coefficients $\mathbf{c}^{1,\delta}$ recovered from the ℓ_1 -minimization problem $(P_{1,\delta})$ do not blow up in the presence of the truncation error δ . We first assume that \mathbf{x} is a fixed point in space and subsequently extend the analysis over the entire spatial domain \mathcal{D} .

Lemma 3.6 (A condition on sparsity for stability of $(P_{1,\delta})$) *Let $u_p^0(\mathbf{y}) := \sum_{\alpha \in \Lambda_{p,d}} c_\alpha^0 \psi_\alpha(\mathbf{y})$ with $c_\alpha^0 = 0$ for $\alpha \notin \Lambda_{p_\epsilon, \nu_\epsilon}$ be the sparse Legendre PC approximation of $u(\mathbf{x}, \mathbf{y})$ at a spatial point \mathbf{x} where the sparse index set $\Lambda_{p_\epsilon, \nu_\epsilon}$ is defined in (33). Assume that the vector of PC coefficients \mathbf{c}^0 satisfies the sparsity condition*

$$\|\mathbf{c}^0\|_0 = |\Lambda_{p_\epsilon, \nu_\epsilon}| < (1 + 1/\mu(\Psi))/4. \quad (49)$$

Let $u_p^{1,\delta}(\mathbf{y}) := \sum_{\alpha \in \Lambda_{p,d}} c_\alpha^{1,\delta} \psi_\alpha(\mathbf{y})$ be the approximation of $u(\mathbf{y})$ with coefficients $\mathbf{c}^{1,\delta}$ computed from the ℓ_1 -minimization problem $(P_{1,\delta})$. Then, with probability at least $1 - \exp(-\frac{N}{8P^4 c_{p,d}})$ (on the N samples $\{u(\mathbf{y}_i)\}_{i=1}^N$) and for all $\delta \geq 0$, the solution $u_p^{1,\delta}$ must obey

$$\mathbb{E} \left[\left(u_p^0 - u_p^{1,\delta} \right)^2 \right] \leq \frac{4}{3N} \frac{(\delta + \|\Psi \mathbf{c}^0 - \mathbf{u}\|_2)^2}{1 - \mu(\Psi) (4\|\mathbf{c}^0\|_0 - 1)}. \quad (50)$$

Proof. Using Theorem 3.1 of [31], we obtain that if \mathbf{c}^0 satisfies the sparsity condition $\|\mathbf{c}^0\|_0 < (1 + 1/\mu(\Psi))/4$, then

$$\sum_{\alpha} \left(c_\alpha^{1,\delta} - c_\alpha^0 \right)^2 \|\psi_\alpha\|_2^2 \leq \frac{(\delta + \|\Psi \mathbf{c}^0 - \mathbf{u}\|_2)^2}{1 - \mu(\Psi) (4\|\mathbf{c}^0\|_0 - 1)}$$

where $\|\psi_\alpha\|_2$ is the ℓ_2 -norm of the column of Ψ corresponding to index α . The presence of $\|\psi_\alpha\|_2$ is due to the fact that the columns of Ψ are not normalized. Next, using McDiarmid's inequality and the independence of the entries $\Psi[i, j]$ for distinct i s, we obtain that

$$\text{Prob} \left[\sum_{\alpha} \left(c_\alpha^{1,\delta} - c_\alpha^0 \right)^2 \|\psi_\alpha\|_2^2 \leq \frac{3N}{4} \sum_{\alpha} \left(c_\alpha^{1,\delta} - c_\alpha^0 \right)^2 \|\psi_\alpha\|_2^2 \right] \leq \exp \left(-\frac{N}{8\|\psi_\alpha\|_{L^\infty(\Gamma)}^4} \right) \quad (51)$$

We conclude using Lemma 3.5 and the fact that due to the orthonormality of $\{\psi_\alpha(\mathbf{y})\}$ we have $\mathbb{E} \left[\left(u_p^0 - u_p^{1,\delta} \right)^2 \right] = \|\mathbf{c}^0 - \mathbf{c}^{1,\delta}\|_2^2$. \square

Remark: The error bound in (50) is not tight; in fact, the actual error is significantly smaller than the upper bound given in (54). More importantly, according to [31], the sparsity condition (49) is unnecessarily too restrictive. In practice, both far milder sparsity conditions are needed and much better actual errors are achieved.

Remark: We will later use the sparsity condition (49) to derive the sufficient condition (22) (together with (24)) on the number N of random samples needed for a successful recovery. As the condition (49) is too restrictive, the theoretical lower bound on N given in (22) and (24) is too pessimistic.

Remark: According to Lemma 3.6, we do not need to know *a priori* the sparse index set $\Lambda_{p_\epsilon, \nu_\epsilon}$; only the sparsity condition (49) is required.

Remark: We stated Lemma 3.6 for the case where the sparsity of the PC expansion is due to the fact that the effective dimensionality ν_ϵ is potentially smaller than d . However, as far as the stability condition (49) is satisfied, similar stability results are valid for situations where dominant basis are defined over all the dimensions.

Remark: With slight modifications, a similar argument as in Lemma 3.6 may be asserted for the solution of ℓ_0 -minimization problem $(P_{0,\delta})$. Specifically, in that case, we only require a sparsity limit $\|\mathbf{c}^0\|_0 = |\Lambda_{p_\epsilon, \nu_\epsilon}| < (1 + 1/\mu(\Psi))/2$ to achieve the error estimate $\mathbb{E} \left[\left(u_p^0 - u_p^{0,\delta} \right)^2 \right] \leq \frac{4}{3N} \frac{(\delta + \|\Psi \mathbf{c}^0 - \mathbf{u}\|_2)^2}{1 - \mu(\Psi)(2\|\mathbf{c}^0\|_0 - 1)}$.

Notice that the normalized truncation error

$$\epsilon_N^2 := \frac{\|\Psi \mathbf{c}^0 - \mathbf{u}\|_2^2}{N} \quad (52)$$

is the sample average estimate of the mean-squares sparse approximation error $\mathbb{E} \left[(u - u_p^0)^2 \right]$ at the point \mathbf{x} and is a function of samples $\{\mathbf{y}\}_{i=1}^N$ in addition to the order p_ϵ and the dimensionality ν_ϵ of the sparse PC expansion. For a given set of N random independent samples $\{\mathbf{y}\}_{i=1}^N$ and $\{u(\mathbf{y})\}_{i=1}^N$, we may bound ϵ_N^2 in probability using McDiarmid's inequality, i.e.,

$$Prob \left[\epsilon_N^2 \geq \mathbb{E} \left[(u - u_p^0)^2 \right] + r \right] \leq \exp \left(-2N \frac{r^2}{\|u - u_p^0\|_{L^\infty(\Gamma)}^4} \right). \quad (53)$$

Although our sparse approximations are point-wise in space, we are ultimately interested in deriving suitable global stability and error estimates for our sparse reconstructions. Such extensions are readily available from Lemma 3.6 and are stated in the following corollary.

Corollary 3.7 *Let $u_p^0(\mathbf{x}, \mathbf{y}) := \sum_{\alpha \in \Lambda_{p,d}} c_\alpha^0(\mathbf{x}) \psi_\alpha(\mathbf{y})$ with $c_\alpha^0(\mathbf{x}) = 0$ for $\alpha \notin \Lambda_{p_\epsilon, \nu_\epsilon}$ be the sparse Legendre PC approximation of $u(\mathbf{x}, \mathbf{y})$ where the sparse index set $\Lambda_{p_\epsilon, \nu_\epsilon}$ is defined in (33). Assume that the vector of PC coefficients $\mathbf{c}^0(\mathbf{x})$ satisfies the sparsity condition $\|\mathbf{c}^0\|_0 = |\Lambda_{p_\epsilon, \nu_\epsilon}| < (1 + 1/\mu(\Psi))/4$. Let $u_p^{1,\delta}(\mathbf{x}, \mathbf{y}) := \sum_{\alpha \in \Lambda_{p,d}} c_\alpha^{1,\delta}(\mathbf{x}) \psi_\alpha(\mathbf{y})$ be the approximation of $u(\mathbf{x}, \mathbf{y})$ with coefficients $\mathbf{c}^{1,\delta}(\mathbf{x})$ computed from the ℓ_1 -minimization problem $(P_{1,\delta})$ at any point $\mathbf{x} \in \mathcal{D}$. Then, with probability at least $1 - \exp(-\frac{N}{8P^{4c_{p,d}}})$ (on the N samples*

$\{u(\mathbf{y}_i)\}_{i=1}^N$) and for all $\delta \geq 0$, the solution $u_p^{1,\delta}$ must obey

$$\|u_p^0 - u_p^{1,\delta}\|_{L^2(\mathcal{D}, L^2(\Gamma))}^2 \leq \frac{4}{3N} \frac{(\delta + \|\Psi \mathbf{c}^0 - \mathbf{u}\|_{L^2(\mathcal{D}, \ell_2(\mathbb{R}^N))})^2}{1 - \mu(\Psi)(4\|\mathbf{c}^0\|_0 - 1)}. \quad (54)$$

Furthermore, the normalized error $\epsilon_N^2 =: \frac{\|\Psi \mathbf{c}^0 - \mathbf{u}\|_{L^2(\mathcal{D}, \ell_2(\mathbb{R}^N))}^2}{N}$ is bounded from above in probability through

$$\text{Prob} \left[\epsilon_N^2 \geq \|u - u_p^0\|_{L^2(\mathcal{D}, L^2(\Gamma))}^2 + r \right] \leq \exp \left(-2N \frac{r^2}{\|u - u_p^0\|_{L^2(\mathcal{D}, L^\infty(\Gamma))}^4} \right). \quad (55)$$

We have now all the necessary tools to proceed with the proof of our main result stated in Theorem 3.3 which is primarily a direct consequence of Lemma 3.2, Theorem 3.4, and Corollary 3.7.

Proof of Theorem 3.3. We first note that, given the conditions of Lemma 3.2, the solution to problem (2) admits a sparse Legendre PC expansion u_p with sparsity $S = |\Lambda_{p_\epsilon, \nu_\epsilon}| \lesssim \epsilon^{-1/\rho}$ when $\epsilon \geq \exp\left(-\left(\frac{d}{c_1}\right)^\kappa\right)$ for some constants c_1 and (arbitrary) $\rho > 0$. Notice that the sparse approximation u_p has an accuracy better than ϵ in the $H_0^1(\mathcal{D}, L^\infty(\Gamma))$ sense. Based on Theorem 3.4, it is sufficient to have random solution samples of size $N \geq 64P^{4c_p, d}(\ln P)S$, to meet the sparsity requirement $S = |\Lambda_{p_\epsilon, \nu_\epsilon}| < (1 + 1/\mu(\Psi))/4$, in Corollary 3.7, with probability at least $1 - 4P^{2-2S_{\max}}$ where $S_{\max} := \frac{N}{64P^{4c_p, d}(\ln P)}$. On the other hand, given N random samples of solution, we require $\epsilon \geq \frac{1}{S_{\max}^\rho}$ to satisfy the sparsity condition. Given Corollary 3.7 and using the triangular and Poincaré inequalities, with probability at least $1 - P^{-8S_{\max}}$, we have

$$\begin{aligned} \|u - u_p^{1,\delta}\|_{L^2(\mathcal{D}, L^2(\Gamma))} &\leq c_{\mathcal{D}} \|u - u_p^0\|_{H_0^1(\mathcal{D}, L^2(\Gamma))} + \|u_p^0 - u_p^{1,\delta}\|_{L^2(\mathcal{D}, L^2(\Gamma))} \\ &\leq c_{\mathcal{D}}\epsilon + \frac{2}{\sqrt{3N}} \frac{\delta + \|\Psi \mathbf{c}^0 - \mathbf{u}\|_{L^2(\mathcal{D}, \ell_2(\mathbb{R}^N))}}{\sqrt{1 - \mu(\Psi)(4S - 1)}}, \end{aligned}$$

for all $\delta > 0$. Moreover, by choosing $r = \frac{1}{4} \|u - u_p^0\|_{L^2(\mathcal{D}, L^\infty(\Gamma))}^2$ in (55),

$$\epsilon_N^2 =: \frac{\|\Psi \mathbf{c}^0 - \mathbf{u}\|_{L^2(\mathcal{D}, \ell_2(\mathbb{R}^N))}^2}{N} \leq \|u - u_p^0\|_{L^2(\mathcal{D}, L^2(\Gamma))}^2 + \frac{1}{4} \|u - u_p^0\|_{L^2(\mathcal{D}, L^\infty(\Gamma))}^2 \leq \frac{5}{4} c_{\mathcal{D}}^2 \epsilon^2$$

with probability at least $1 - P^{-8S_{\max}P^{4c_p, d}}$. Finally, by taking

$$c_4 := c_{\mathcal{D}} \left(1 + \frac{\sqrt{5}}{\sqrt{3}\sqrt{1 - \mu(\Psi)(4S - 1)}} \right) \quad \text{and} \quad c_5 := \frac{2}{\sqrt{3}\sqrt{1 - \mu(\Psi)(4S - 1)}},$$

we arrive at the statement of the Theorem 3.3. \square

3.6 Choosing the truncation error tolerance δ

An important component of the sparse approximation using $(P_{1,\delta})$ and $(P_{0,\delta})$ is the selection of the truncation error tolerance δ . Although the stability bounds given in Lemma 3.6 and Corollary 3.7 are valid for any $\delta \geq 0$, the actual error and the sparsity level of the solution to $(P_{1,\delta})$ and $(P_{0,\delta})$ depend on the choice of δ . Ideally, we desire to choose $\delta \approx \|\Psi \mathbf{c}^0 - \mathbf{u}\|_2$; while larger values of δ deteriorate the accuracy of the approximation, as in Lemma 3.6, smaller choices of δ may result in over-fitting the solution samples and, thus, less sparse solutions. In practice, as the exact values of the PC coefficients \mathbf{c}^0 are not known, the exact values of the truncation error $\|\Psi \mathbf{c}^0 - \mathbf{u}\|$ and, consequently, δ are not known *a priori*. Therefore, δ has to be estimated, for instance, using statistical techniques such as the cross-validation [11, 59].

In this work, we propose a heuristic cross-validation algorithm to estimate δ . We first divide the N available solution samples to N_r reconstruction and N_v validation samples such that $N = N_r + N_v$. The idea is to repeat the solution of $(P_{1,\delta})$ (or $(P_{0,\delta})$) on the reconstruction samples and with multiple values of truncation error tolerance δ_r . We then set $\delta = \sqrt{\frac{N}{N_r}} \hat{\delta}_r$ in which $\hat{\delta}_r$ is such that the corresponding truncation error on the N_v validation samples is minimum. This is simply motivated by the fact that the truncation error on the validation samples is large for values of δ_r considerably larger and smaller than $\|\Psi \mathbf{c}^0 - \mathbf{u}\|_2$ evaluated using the reconstruction samples. While the former is expected from the upper bound on the approximation error in Lemma 3.6, the latter is due to the over-fitting the reconstruction samples. The following exhibit outlines the estimation of δ using the above cross-validation approach:

Algorithm for cross-validation estimation of δ :

- Divide the N solution samples to N_r reconstruction and N_v validation samples.
- Choose multiple values for δ_r such that the exact truncation error $\|\Psi \mathbf{c}^0 - \mathbf{u}\|_2$ of the reconstruction samples is within the range of δ_r values.
- For each value of δ_r , solve $(P_{1,\delta})$ (or $(P_{0,\delta})$) on the N_r reconstruction samples.
- For each value of δ_r , compute the truncation error $\delta_v := \|\Psi \mathbf{c}^{1,\delta_r} - \mathbf{u}\|_2$ (or $\delta_v := \|\Psi \mathbf{c}^{0,\delta_r} - \mathbf{u}\|_2$) of the N_v validation samples.
- Find the minimum value of δ_v and its corresponding $\hat{\delta}_r := \delta_r$.
- Set $\delta = \sqrt{\frac{N}{N_r}} \hat{\delta}_r$.

In the numerical experiments of Section 4, we repeat the above cross-validation algorithm for multiple replications of the reconstruction and validation samples. The estimate of $\delta = \sqrt{\frac{N}{N_r}} \hat{\delta}_r$ is then based on the value of $\hat{\delta}_r$ for which the average of the corresponding truncation errors δ_v , over all replications of the validation samples, is minimum. This resulted in more accurate solutions in our numerical experiments.

3.7 Algorithms

There are several numerical algorithms for solving problems $(P_{0,\delta})$ and $(P_{1,\delta})$ each with different optimization kernel, computational complexity, and degree of accuracy. An in-depth discussion on the performance of these algorithms is outside the scope of the present work; however, below we name some of the available options for each problem and briefly describe the algorithms that have been utilized in our numerical experiments. For comprehensive discussions on this subject, the interested reader is referred to [13, 57, 35, 4, 63, 56].

Problem $(P_{0,\delta})$: A brute force search through all possible support sets in order to identify the correct sparsity for the solution $\mathbf{c}^{0,\delta}$ of $(P_{0,\delta})$ is NP-hard and not practical. Greedy pursuit algorithms form a major class of schemes to tackle the solution of $(P_{0,\delta})$ with a tractable computational cost. Instead of performing an exhaustive search for the support of the sparse solution, these solvers successively find one or more components of the solution that result in the largest improvement in the approximation. Some of the standard greedy pursuit algorithms are Orthogonal Marching Pursuit (OMP) [51, 26], Regularized OMP (ROMP)[45], Stagewise OMP (StOMP) [30], Compressive Sampling MP (CoSaMP) [44], Subspace Pursuit [23], and Iterative Hard Thresholding (IHT) [25]. Under well-defined conditions, all of the above schemes provide stable and accurate solutions to $(P_{0,\delta})$ in a reasonable time.

In the present study, we employ the OMP algorithm to approximate the solution of $(P_{0,\delta})$. Starting from $\mathbf{c}^{0,\delta,(0)} = \mathbf{0}$ and an empty active column set of Ψ , at any iteration k , OMP identifies only one column to be added to the active column set. The column is chosen such that the ℓ_2 -norm of the residual, $\|\Psi \mathbf{c}^{0,\delta,(k)} - \mathbf{u}\|_2$, is maximally reduced. Having specified the active column set, a least-squares problem is solved to compute the solution $\mathbf{c}^{0,\delta,(k)}$. The iterations are continued until the error truncation tolerance δ is achieved. In general, the complexity of the OMP algorithm is $\mathcal{O}(S \cdot N \cdot P)$ where $S := \|\mathbf{c}^{0,\delta}\|_0$ is number of non-zero (dominant) entries of $\mathbf{c}^{0,\delta}$. The following exhibit depicts an step-by-step implementation of the OMP algorithm.

Orthogonal Matching Pursuit (OMP) Algorithm:

- Set $k = 0$.
 - Set the initial solution $\mathbf{c}^{0,\delta,(0)} = \mathbf{0}$ and residual $\mathbf{r}^{(0)} = \mathbf{u} - \Psi\mathbf{c}^{0,\delta,(0)} = \mathbf{u}$.
 - Set the solution support index set $\mathcal{I}^{(0)} = \emptyset$.
- While $\|\mathbf{u} - \Psi\mathbf{c}^{0,\delta,(k)}\|_2 > \delta$ perform:
 - For all $j \notin \mathcal{I}^{(k)}$ evaluate $\epsilon(j) = \|\boldsymbol{\psi}_j\alpha_j - \mathbf{r}^{(k)}\|_2$ with $\alpha_j = \boldsymbol{\psi}_j^T \mathbf{r}^{(k)} / \|\boldsymbol{\psi}_j\|_2^2$.
 - Set $k = k + 1$.
 - Update the support index set $\mathcal{I}^{(k)} = \mathcal{I}^{(k-1)} \cup \{\arg \min_j \epsilon(j)\}$.
 - Solve for $\mathbf{c}^{0,\delta,(k)} = \arg \min_{\mathbf{c}^{0,\delta}} \|\mathbf{u} - \Psi\mathbf{c}^{0,\delta}\|_2$ subject to $\text{Support}\{\mathbf{c}^{0,\delta}\} = \mathcal{I}^{(k)}$.
 - Update the residual $\mathbf{r}^{(k)} = \mathbf{u} - \Psi\mathbf{c}^{0,\delta,(k)}$
- Output the solution $\mathbf{c}^{0,\delta} = \mathbf{c}^{0,\delta,(k)}$.

Although we chose OMP in our analysis, we note that further studies are needed to identify the most appropriate greedy algorithm for the purpose of this study.

Problem ($P_{1,\delta}$): The majority of available solvers for ℓ_1 -minimization are based on alternative formulations of ($P_{1,\delta}$), such as the ℓ_1 -norm regularized least-squares problem

$$(QP_\lambda) : \quad \min_{\mathbf{c}} \frac{1}{2} \|\Psi\mathbf{c} - \mathbf{u}\|_2^2 + \lambda \|\mathbf{W}\mathbf{c}\|_1, \quad (56)$$

or the LASSO problem, [54],

$$(LS_\tau) : \quad \min_{\mathbf{c}} \frac{1}{2} \|\Psi\mathbf{c} - \mathbf{u}\|_2^2 \quad \text{subject to} \quad \|\mathbf{W}\mathbf{c}\|_1 \leq \tau. \quad (57)$$

It can be shown that for an appropriate choice of scalars δ , λ , and τ , the problems ($P_{1,\delta}$), (QP_λ), and (LS_τ) share the same solution [57, 13, 56]. Among others, the problem (QP_λ) is of particular interest as it is an unconstrained optimization problem. Numerous solvers based on the *active set* [50, 34], *interior-point continuation* [20, 40] and *projected gradient* [24, 22, 37, 12, 9, 57, 3, 4] methods have been developed for solving the above formulations of the ℓ_1 -minimization problem.

In our numerical experiments, we adopt the Spectral Projected Gradient algorithm (SPGL1) proposed in [57] and implemented in the MATLAB package SPGL1 [5] to solve the ℓ_1 -minimization problem ($P_{1,\delta}$) in (21). SPGL1 is based on exploring the so-called Pareto curve, describing the tradeoff between the ℓ_2 -norm of the truncation error $\|\Psi\mathbf{c} - \mathbf{u}\|_2$ and the (weighted) ℓ_1 -norm of the solution $\|\mathbf{W}\mathbf{c}\|_1$, for successive solution iterations. At each iteration, the

LASSO problem (57) is solved using the spectral projected gradient technique with a worst-case complexity of $\mathcal{O}(P \ln P)$ where P is the number of columns in $\mathbf{\Psi}$. Given the error tolerance δ , a scalar equation is solved to identify a value for τ such that the (LS_τ) solution of (57) is identical to that of $(P_{1,\delta})$ in (21). Besides being efficient for large-scale systems where $\mathbf{\Psi}$ may not be available explicitly, the SPGL1 algorithm is specifically effective for our application of interest as the truncation error $\|\mathbf{\Psi}\mathbf{c} - \mathbf{u}\|_2$ is known only approximately.

In the next section, we explore some aspects of the proposed scheme through its application to a 1- D (in space) elliptic stochastic PDE with high-dimensional random diffusion coefficients.

4 Numerical examples

We consider the solution of a one-dimensional, i.e., $D = 1$, version of problem (2),

$$\begin{aligned} -\frac{d}{dx} \left(a(x, \omega) \frac{du(x, \omega)}{dx} \right) &= 1, \quad x \in \mathcal{D} = (0, 1), \\ u(0, \omega) &= u(1, \omega) = 0, \end{aligned} \quad (58)$$

where the stochastic diffusion coefficient $a(x, \omega)$ is given by the expansion

$$a(x, \omega) = \bar{a} + \sigma_a \sum_{i=1}^d \sqrt{\lambda_i} \phi_i(x) y_i(\omega). \quad (59)$$

Here, $\{\lambda_i\}_{i=1}^d$ and $\{\phi_i(x)\}_{i=1}^d$ are, respectively, d largest eigenvalues and the corresponding eigenfunctions of the Gaussian covariance kernel

$$C_{aa}(x_1, x_2) = \exp \left[-\frac{(x_1 - x_2)^2}{l_c^2} \right], \quad (60)$$

in which l_c is the correlation length of $a(x, \omega)$ that prescribes the decay of the spectrum of C_{aa} in (60). Random variables $\{y_i(\omega)\}_{i=1}^d$ are assumed to be independent and uniformly distributed on $[-1, 1]$. The coefficient σ_a controls the variability of $a(x, \omega)$.

We verify the accuracy and efficiency of the present sparse approximation schemes for both moderate and high-dimensional diffusion coefficient $a(x, \omega)$. These two cases are obtained, respectively, by assuming $(l_c, d) = (1/5, 14)$ and $(l_c, d) = (1/14, 40)$ in (60) and (59). We further assume that $\bar{a} = 0.1$, $\sigma_a = 0.03$ when $d = 14$, and $\sigma_a = 0.021$ when $d = 40$. These choices ensure

that all realizations of $a(x, \omega)$ are strictly positive on $\mathcal{D} = (0, 1)$. Table 1 summarizes the assumed parameters for the two test cases.

Table 1

Choices of parameters defining the stochastic description of diffusion coefficient $a(x, \omega)$ in Eq. (59).

Case	\bar{a}	σ_a	l_c	d
I	0.1	0.030	1/5	14
II	0.1	0.021	1/14	40

For both cases, the spatial discretization is done by the Finite Element Method using quadratic elements. A mesh convergence analysis is performed to ensure that spatial discretization errors are inconsequential.

The solution statistics are computed using the conventional Monte Carlo simulation, the isotropic sparse grid stochastic collocation with the Clenshaw-Curtis abscissas [60, 1], and the proposed sparse approximation techniques. We use the Basis Pursuit Denoising (BPDN) solver implemented in `SPGL1` [5, 57] to solve the ℓ_1 -minimization problem $(P_{1,\delta})$ and the Orthogonal Matching Pursuit (OMP) solver in `SparseLab` [28] to approximate the ℓ_0 -minimization problem $(P_{0,\delta})$. We compare the errors in the mean, standard deviation, and root mean-squares of the solution error at $x = 0.5$ using the above methods. The details of the analysis are reported below.

4.1 Case I: $d=14$

We consider an increasing number $N = \{29, 120, 200, 280, 360, 421, 600\}$ of random solution samples to evaluate the solution u at $x = 0.5$ and, consequently, to compute the PC coefficients of the solution using ℓ_1 - and ℓ_0 -minimization. These samples are nested in the sense that we recycle the previous samples when we perform calculations with larger sample sizes. The nested sampling property of our scheme is of paramount importance in large scale calculations where the computational cost of each solution evaluation is enormous. We note that sample sizes $N = 29$ and $N = 421$, respectively, correspond to the number of nested abscissas in the level $l = 1$ and level $l = 2$ of the isotropic stochastic collocation with the Clenshaw-Curtis rule.

As elucidated in Section 3.5, the accuracy of our sparse reconstruction depends on the mutual coherence $\mu(\Psi)$, the sample size N , and the truncation error $\|\Psi \mathbf{c} - \mathbf{u}\|_2$ (hence δ). In order to reduce the approximation error, we need

to reduce $\|\Psi\mathbf{c} - \mathbf{u}\|_2$, which may be done by increasing p and, therefore, P . However, with a fixed number N of samples, an increase in P may result in a larger mutual coherence and, thus, the degradation of the reconstruction accuracy. Therefore, in practice, we start by approximating the lower order PC expansions when N is small and increase p when larger number of samples become available. Notice that such an adaptivity with respect to the order p is a natural way of refining the accuracy of PC expansions, for instance, when the intrusive stochastic Galerkin scheme is adopted [36]. In particular, in this example, for sample sizes $N = \{29, 120\}$, we attempt to estimate the coefficients of the 3rd-order Legendre PC expansion, i.e. $p = 3$ and $P = 680$. For larger sample sizes N , we also include the first 320 basis function from the 4th-order chaos, thus resulting in $P = 1000$. Since all of the 4th-order basis functions are not employed, we need to describe the ordering of our basis construction. We sort the elements of $\{\psi_{\alpha}(\mathbf{y})\}$ such that, for any given order p , the random variables y_i with smaller indices i contribute first in the basis.

For each analysis, we estimate the truncation error tolerance δ based on the cross-validation algorithm described in Section 3.6. For each N , we use $N_r \approx 3N/4$ of the samples (reconstruction set) to compute the PC coefficients \mathbf{c}^{1,δ_r} and the rest of the samples (validation set) are used to evaluate the truncation error δ_v . The cross-validation is performed for four replications of reconstruction and validation sample sets. We then find the value $\hat{\delta}_r$ that minimizes the average of δ_v over the four replications of the cross-validation samples. Given an estimate of the truncation error tolerance $\delta \approx \sqrt{4/3}\hat{\delta}_r$, we then use all N samples to compute the coefficients $\mathbf{c}^{1,\delta}$.

Figure 2 compares the ‘exact’ PC coefficients with those obtained using BPDN and OMP solvers. We only demonstrate the results corresponding to sample sizes $N = \{120, 600\}$. An ‘exact’ solution is computed using the level $l = 8$ stochastic collocation for which the approximation errors are negligible in our comparisons. We observe that BPDN tends to give less sparse solutions compared to OMP. This is due to the facts that *i*) the solution is not exactly sparse, i.e., there are many non-zero (although negligible) coefficients c_{α} , *ii*) the ℓ_1 cost function does not impose a sufficiently large penalty on the small coefficients as does the ℓ_0 cost function, and *iii*) the truncation error tolerance δ may be under-estimated. To reduce this issue, a number of modifications, including the reweighted ℓ_1 -minimization [18, 62, 43, 39], have been introduced in the literature that are the subjects of our future work. In contrary, OMP results in more sparse solutions as it adds basis function one-at-a-time until the residual falls below the truncation error. However, as is seen in Figs. 2 (b) and (d), a number of small coefficients are still over-estimated. This is primarily due to under-estimation of the truncation error tolerance δ in the cross-validation algorithm.

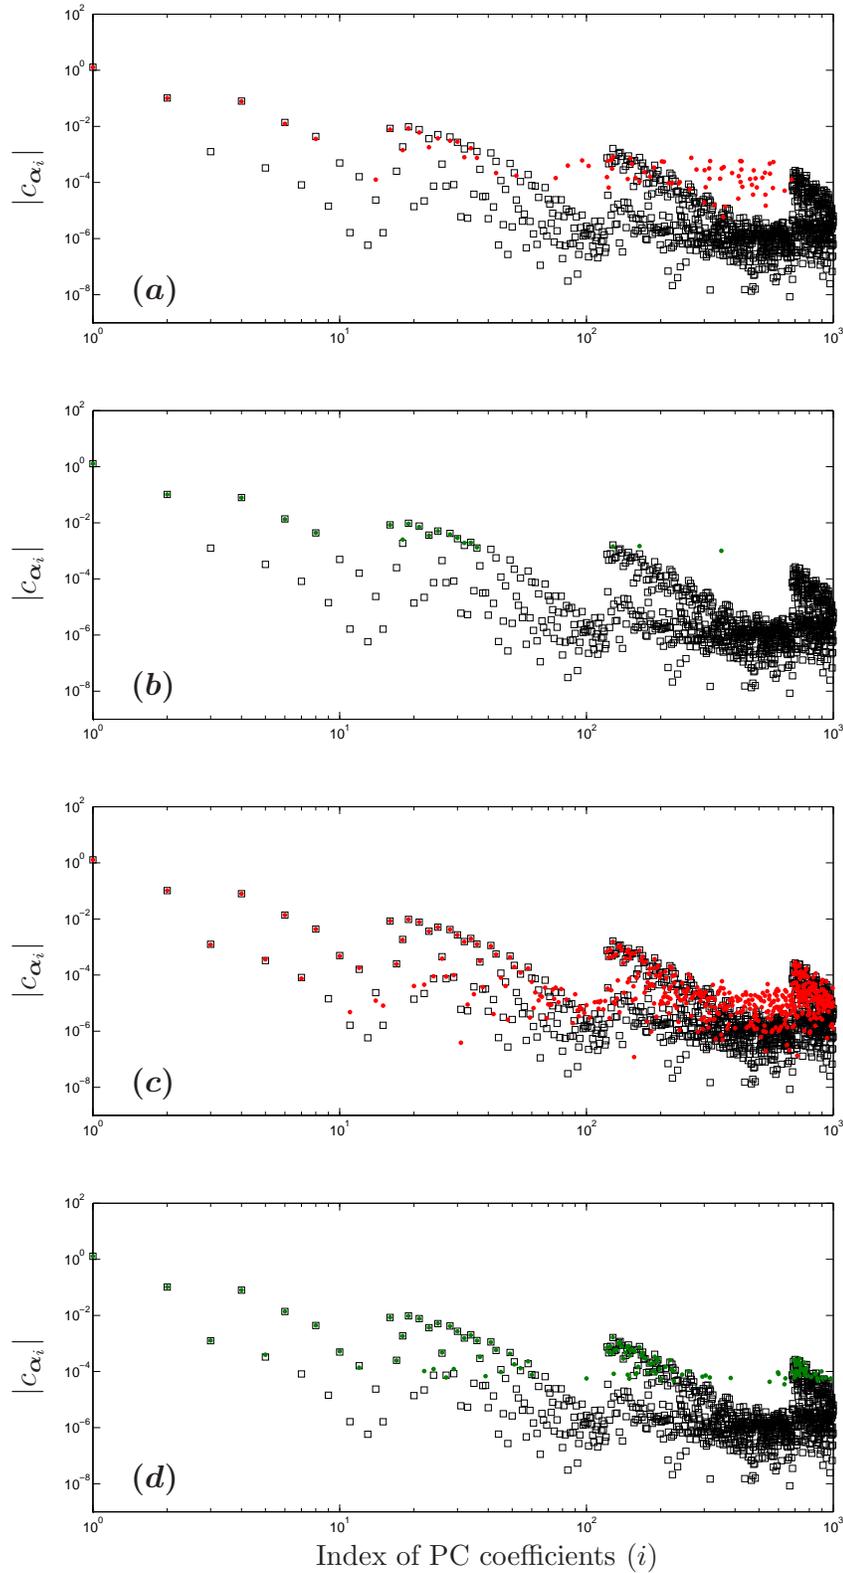


Fig. 2. Approximation of polynomial chaos (PC) coefficients \mathbf{c} of $u(0.5, \mathbf{y})$ using BPDN and OMP for $d = 14$. (a) BPDN with $N = 120$ samples, (b) OMP with $N = 120$ samples, (c) BPDN with $N = 600$ samples, and (d) OMP with $N = 600$ samples. ‘Exact’ coefficients computed from level 8 stochastic collocation with the Clenshaw-Curtis abscissas (\square); BPDN and OMP (\bullet).

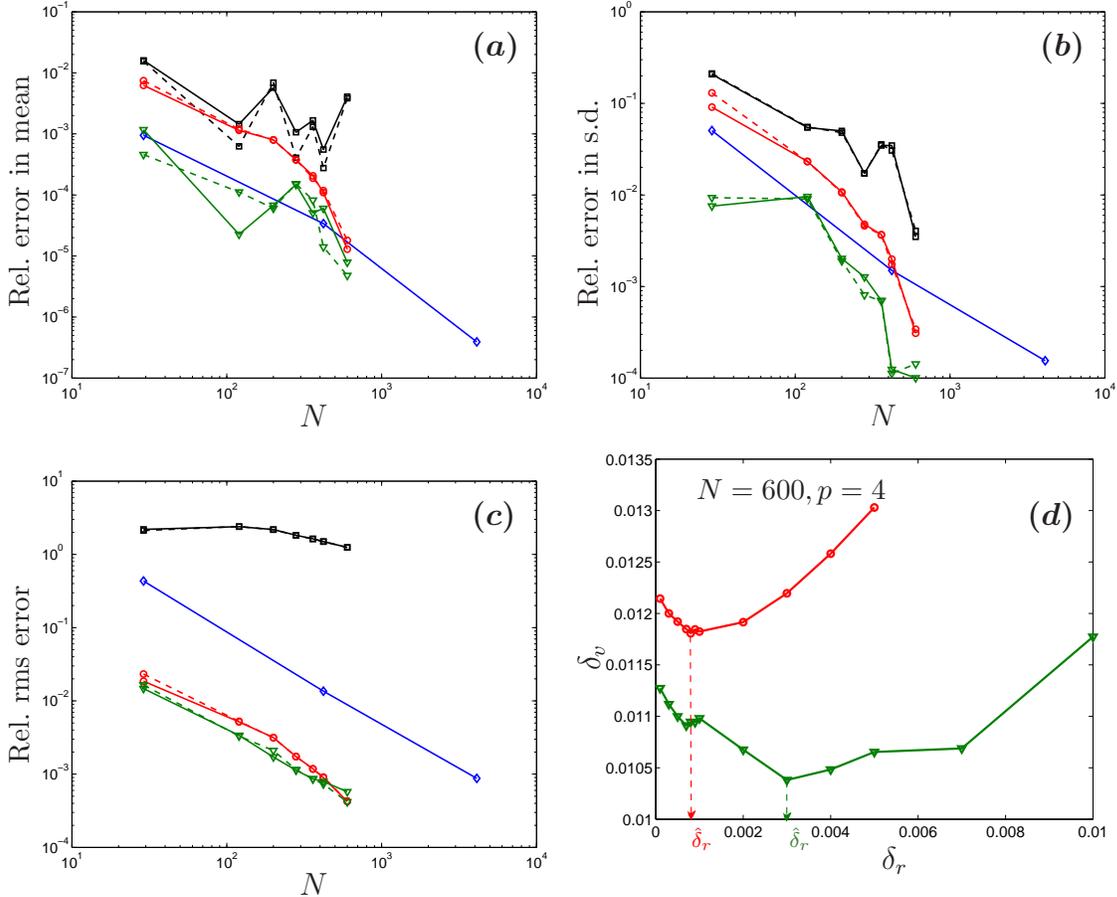


Fig. 3. Comparison of relative error in solution statistics at $x = 0.5$ for the Monte Carlo simulation, isotropic sparse grid stochastic collocation with the Clenshaw-Curtis abscissas, and the proposed sparse approximations (BPDN and OMP) for $d = 14$. Two sets of independent random samples of $u(0.5, \mathbf{y})$ are generated first and are used for the Monte Carlo simulation, BPDN, and OMP. The solid and dashed lines correspond to the first and second sets of samples, respectively. (a) Relative error in the mean; (b) Relative error in the standard deviation; (c) Relative root-mean-squares (rms) error; and (d) Estimation of δ using cross-validation: δ is computed from δ_r for which δ_v is minimum. (Monte Carlo simulation (\square); stochastic collocation (\diamond); BPDN (\circ); OMP (∇)).

The convergence of the mean, standard deviation, and root mean-squares of the approximation error for $u(0.5, \mathbf{y})$ is illustrated in Figs. 3 (a), (b), and (c), respectively. For the case of stochastic collocation, we apply sparse grid quadrature (cubature) integration rule to directly compute the mean and the standard deviation. The root mean-squares error of the Monte Carlo and the stochastic collocation solution are evaluated by estimating the corresponding PC coefficients using sampling and sparse grid quadrature integration, respectively, and then comparing them with the exact coefficients.

To make a meaningful comparison, for each N , the samples used to compute the solution statistics by the conventional Monte Carlo, BPDN, and OMP are identical. In this sense, the sparse approximation using ℓ_1 - and ℓ_0 -minimizations may be viewed as only post-processing steps in the Monte Carlo simulation. As the sample sizes are finite, the estimates of the PC coefficients \mathbf{c} are sample dependent and are in fact random. To demonstrate the convergence of the algorithm with respect to different sets of samples, we repeat the analysis for two independent sets of N samples and report the corresponding statistics errors with solid and dashed lines. Although for different solution samples of size N the estimates of \mathbf{c} and the solution statistics are not identical, the approximation converges, with large probability, for any set of samples with sufficiently large size N (see Theorem 3.3).

Figure 3 (d) illustrates the statistical estimation of δ using the cross-validation approach described in Section 3.6. The estimation of δ is slightly different in BPDN and OPM, this is a consequence of different reconstruction accuracy of these techniques. Moreover, the solution of the BPDN algorithm is less sensitive to small perturbations in the truncation error δ compared to that of the OMP algorithm. This is justified by the fact that the ℓ_0 -norm is highly discontinuous.

Remark: Despite the conventional implementation of the stochastic collocation approach where the approximation refinement requires a certain number of extra samples, the ℓ_1 - and ℓ_0 -minimizations may be implemented using arbitrary numbers of additional samples, which is an advantage, particularly, when only a limited number of samples is afforded.

4.2 Case II: $d=40$

The objective of this example is to highlight that a sparse reconstruction may lead to significant computational savings for problems with high-dimensional random inputs. Similar to the analysis of Case I described in Section 4.1, we compute the solution statistics using multiple numbers of independent samples. More specifically, we evaluate the solution at $x = 0.5$ for independent samples of size $N = \{81, 200, 400, 600, 800, 1000\}$. The number of grid points in the level $l = 1$ and $l = 2$ of the Clenshaw-Curtis rule in dimension $d = 40$ is $N = 81$ and $N = 3281$, respectively. To obtain a reference solution, the 3rd order PC coefficients \mathbf{c} of the solution at $x = 0.5$ are computed using level $l = 5$ stochastic collocation with the Clenshaw-Curtis rule.

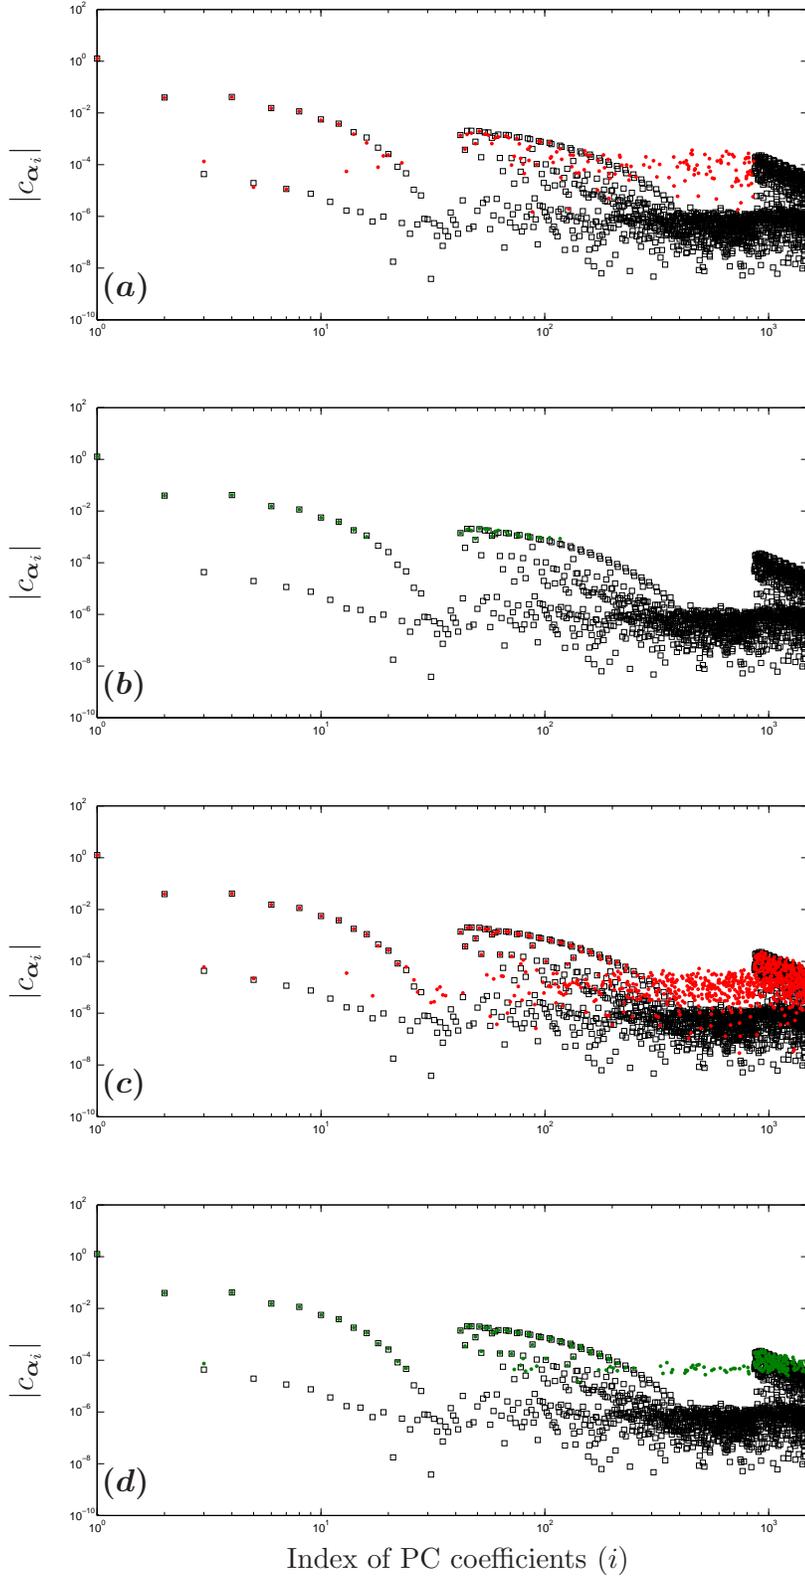


Fig. 4. Approximation of polynomial chaos (PC) coefficients \mathbf{c} of $u(0.5, \mathbf{y})$ using BPDN and OMP for $d = 40$. (a) BPDN with $N = 200$ samples, (b) OMP with $N = 200$ samples, (c) BPDN with $N = 1000$ samples, and (d) OMP with $N = 1000$ samples. ‘Exact’ coefficients computed from level 5 stochastic collocation with the Clenshaw-Curtis abscissas (\square); BPDN and OMP (\bullet).

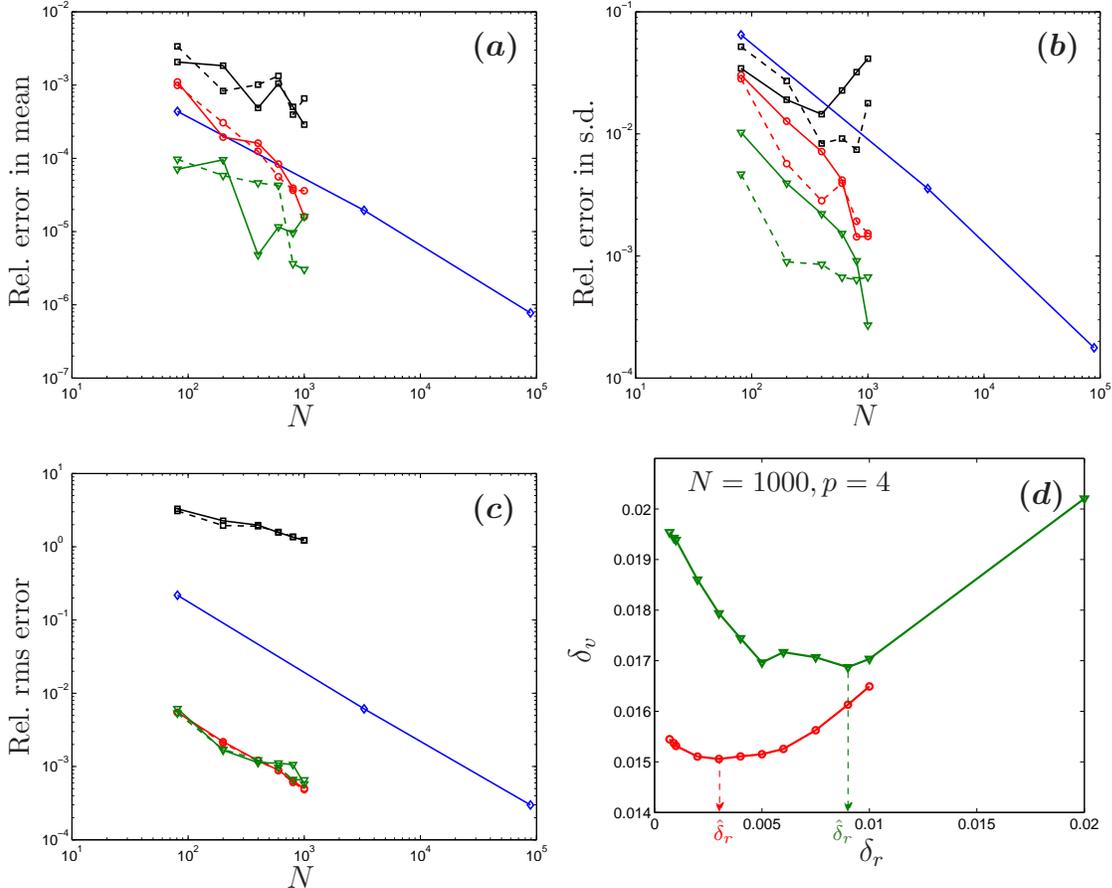


Fig. 5. Comparison of relative error in solution statistics at $x = 0.5$ for the Monte Carlo simulation, isotropic sparse grid stochastic collocation with the Clenshaw-Curtis abscissas, and the proposed sparse approximations (BPDN and OMP) for $d = 40$. Two sets of independent random samples of $u(0.5, \mathbf{y})$ are generated first and are used for the Monte Carlo simulation, BPDN, and OMP. The solid and dashed lines correspond to the first and second sets of samples, respectively. (a) Relative error in mean; (b) Relative error in standard deviation; (c) Relative root-mean-squares (rms) error; and (d) Estimation of δ using cross-validation: δ is computed from δ_r for which δ_v is minimum. (Monte Carlo simulation (\square); stochastic collocation (\diamond); BPDN (\circ); OMP (∇)).

For $N = \{81, 200\}$ we only estimate the coefficients associated with the 2nd-order PC expansion, i.e. $p = 2$ and $P = 861$. For larger sample sizes, we also include the first 639 basis functions from the 3th-order chaos, thus leading to $P = 1500$. For each combination of N and p , we estimate the truncation error δ using an identical cross-validation procedure described in Section 4.2. Figure 4 illustrates the estimation of PC coefficients of $u(0.5, \mathbf{y})$ with BPDN and OMP algorithms with $N = 200$ and $N = 1000$. We again note that the recovered solution from the BPDN algorithm is less sparse as compared to that of the OMP approach, although the over-estimated coefficients (mostly from the second order term) are indeed small. As the samples size N is increased, we are naturally able to recover more dominant coefficients on the expansion.

Figure 5 depicts the convergence of the statistics of the solution as functions of the sample size N as well as one instance of the estimation of the truncation error tolerance δ . The implementation details are similar to those described in Section 4.1 for the case of $d = 14$.

It is worth highlighting that the computational saving of the present sparse approximations (in terms of the number of samples needed to achieve a certain accuracy) compared to the isotropic sparse grid collocation is even larger for the higher-dimensional ($d = 40$) problem. This is due to the fact that the number of samples needed to recover the solution using ℓ_1 - and ℓ_0 -minimization is dictated more by the number of dominant terms in the PC expansion compared to the total number of terms P , as in Theorem 3.1.

5 Conclusion

The present study proposes a *non-intrusive* and *non-adapted* approach based on the *compressive sampling* formalism for the approximation of sparse solution of stochastic PDEs. When sufficiently sparse in the polynomial chaos (PC) basis, the compressive sampling enables an accurate recovery of the solution using a number of random solution samples that is significantly smaller than the cardinality of the PC basis. Sparse PC approximations based on ℓ_0 - and ℓ_1 -minimization approaches have been introduced and implemented using the Basis Pursuit Denoising (BPDN) and the Orthogonal Matching Pursuit (OMP) algorithms, respectively. Probabilistic bounds based on the *concentration of measure* phenomenon have been derived to verify the convergence and the stability of the present sparse constructions. The performance and efficiency of the proposed techniques are explored through their application to a linear elliptic PDE with high-dimensional random diffusion coefficients where the sparsity of the solution with respect to the PC basis is guaranteed. The proposed formalism to recover the sparse PC expansion of stochastic functions is not restricted to the case of the elliptic PDEs, as its underlying applicability assumptions are universal. Although the discussions of this work have been focused on the particular case of the Legendre PC expansions, the proposed framework can be readily extended to other bases such as the Hermite PC (when the random variables \mathbf{y} are standard Gaussian).

Acknowledgments

The first author acknowledges the support of the United States Department of Energy under Stanford's Predictive Science Academic Alliance Program

(PSAAP) for the preliminary stages of his work. The second author acknowledges the support of the National Science Foundation via NSF grant CMMI-092600 and of the United States Department of Energy under Caltech's Predictive Science Academic Alliance Program (PSAAP).

References

- [1] I. Babuška, F. Nobile, and R. Tempone. A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM J. Numer. Anal.*, 45(3):1005–1034, 2007.
- [2] I. Babuška, 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.
- [3] A. Beck and M. Teboulle. A fast iterative shrinkage-threshold algorithm for linear inverse problems. *SIAM J. Imaging Sciences*, 2:183–202, 2009.
- [4] S. Becker, J. Bobin, and E. J. Candes. NESTA: A fast and accurate first-order method for sparse recovery. *ArXiv e-prints*, 2009. Available from <http://arxiv.org/abs/0904.3367>.
- [5] E. van den Berg and M. P. Friedlander. SPGL1: A solver for large-scale sparse reconstruction, June 2007. Available from <http://www.cs.ubc.ca/labs/scl/spgl1>.
- [6] M. Bieri. A sparse composite collocation finite element method for elliptic spdes. Technical Report Research Report No. 2009-08, Seminar für Angewandte Mathematik, SAM, Zürich, Switzerland, 2009.
- [7] M. Bieri, R. Andreev, and C. Schwab. Sparse tensor discretization of elliptic spdes. Technical Report Research Report No. 2009-07, Seminar für Angewandte Mathematik, SAM, Zürich, Switzerland, 2009.
- [8] M. Bieri and C. Schwab. Sparse high order FEM for elliptic sPDEs. *Computer Methods in Applied Mechanics and Engineering*, 198:1149 – 1170, 2009.
- [9] J.M. Bioucas-Dias and M.A.T. Figueiredo. A new TwIST: Two-step iterative shrinking/thresholding algorithms for image restoration. *IEEE Trans. Image Processing*, 16(12):2992–3004, 2007.
- [10] G. Blatman and B. Sudret. An adaptive algorithm to build up sparse polynomial chaos expansions for stochastic finite element analysis. *Probabilistic Engineering Mechanics*, 25(2):183–197, 2010.
- [11] P. Boufounos, M.F. Duarte, and R.G. Baraniuk. Sparse signal reconstruction from noisy compressive measurements using cross validation. In *SSP '07: Proceedings of the 2007 IEEE/SP 14th Workshop on Statistical Signal Processing*, pages 299–303. IEEE Computer Society, 2007.
- [12] K. Bredies and D.A. Lorenz. Linear convergence of iterative soft- thresholding. *SIAM J. Sci. Comp.*, 30(2):657–683, 2008.
- [13] A.M. Bruckstein, D.L. Donoho, and M. Elad. From sparse solutions of

- systems of equations to sparse modeling of signals and images. *SIAM Review*, 51(1):34–81, 2009.
- [14] E.J. Candes and J. Romberg. Quantitative robust uncertainty principles and optimally sparse decompositions. *Found. Comput. Math.*, 6(2):227–254, 2006.
- [15] E.J. Candes and J. Romberg. Sparsity and incoherence in compressive sampling. *Inverse Problems*, 23(3):969–985, 2007.
- [16] E.J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006.
- [17] E.J. Candes and T. Tao. Near optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on information theory*, 52(12):5406–5425, 2006.
- [18] E.J. Candes, M.B. Wakin, and S. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14(5):877–905, 2008.
- [19] S.S. Chen, D.L. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.*, 20:33–61, 1998.
- [20] S.S. Chen, D.L. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM Rev.*, 43(1):129–159, 2001.
- [21] A. Cohen, W. Dahmen, and R. DeVore. Compressed sensing and best k -term approximation. *J. Amer. Math. Soc.*, 22:211–231, 2009.
- [22] P.L. Combettes and V.R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling and Simulation*, 4(4):1168–1200, 2005.
- [23] W. Dai and O. Milenkovic. Subspace pursuit for compressive sensing signal reconstruction. *IEEE Trans. Inf. Theor.*, 55(5):2230–2249, 2009.
- [24] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004.
- [25] T. Blumensath M.E. Davies. Iterative hard thresholding for compressed sensing. *Appl. Comp. Harm. Anal*, 27(3):265–274, 2009.
- [26] G. Davis, S. Mallat, and M. Avellaneda. Adaptive greedy approximation. *J. Constr. Approx.*, 13:57–98, 1997.
- [27] M.K. Deb, I. Babuska, and J.T. Oden. Solution of stochastic partial differential equations using Galerkin finite element techniques. *Comput. Methods Appl. Mech. Engrg.*, 190:6359–6372, 2001.
- [28] D. L. Donoho, I. Drori, V.C. Stodden, and Y. Tsaig. Available from <http://www-stat.stanford.edu/~sparselab/>.
- [29] D.L. Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
- [30] D.L. Donoho, I. Drori, Y. Tsaig, and J.L. Starck. Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit. Technical report, 2006.
- [31] D.L. Donoho, M. Elad, and V.N. Temlyakov. Stable recovery of sparse

- overcomplete representations in the presence of noise. *IEEE Transactions on information theory*, 52(1):6–18, 2006.
- [32] A. Doostan, R. Ghanem, and J. Red-Horse. Stochastic model reduction for chaos representations. *Computer Methods in Applied Mechanics and Engineering*, 196(37-40):3951–3966, 2007.
- [33] A. Doostan and G. Iaccarino. A least-squares approximation of partial differential equations with high-dimensional random inputs. *Journal of Computational Physics*, 228(12):4332–4345, 2009.
- [34] B. Efron, T. Hastie, L. Johnstone, and R. Tibshirani. Least angle regression. *Annals of Statistics*, 32:407–499, 2004.
- [35] M.A.T Figueiredo, R.D. Nowak, and S.J. Wright. Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems. *Selected Topics in Signal Processing, IEEE Journal of*, 1(4):586–597, 2007.
- [36] R. Ghanem and P. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Dover, 2002.
- [37] E.T. Hale, W. Yin, and Y. Zhang. Fixed-point continuation for ℓ_1 -minimization: Methodology and convergence. *SIAM J. on Optimization*, 19(3):1107–1130, 2008.
- [38] S. Hosder, R.W. Walters, and R. Perez. A non-intrusive polynomial chaos method for uncertainty propagation in CFD simulations. In *44th AIAA aerospace sciences meeting and exhibit, AIAA-2006-891*, Reno (NV), 2006.
- [39] M.A. Khajehnejad, W. Xu, A.S. Avestimehr, and B. Hassibi. Improved sparse recovery thresholds with two-step reweighted ℓ_1 minimization. *ArXiv e-prints*, 2010. Available from <http://arxiv.org/abs/1004.0402>.
- [40] S.-J Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky. An interior-point method for large-scale ℓ_1 -regularized least squares. *Selected Topics in Signal Processing, IEEE Journal of*, 1(4):606–617, 2007.
- [41] 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.
- [42] L. Mathelin and M.Y. Hussaini. A stochastic collocation algorithm for uncertainty analysis. Technical Report NAS 1.26:212153; NASA/CR-2003-212153, NASA Langley Research Center, 2003.
- [43] D. Needell. Noisy signal recovery via iterative reweighted ℓ_1 -minimization. In *Proc. Asilomar Conf. on Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 2009.
- [44] D. Needell and J.A. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26(3):301–321, 2008.
- [45] D. Needell and R. Vershynin. Signal recovery from incomplete and inaccurate measurements via Regularized Orthogonal Matching Pursuit. *ArXiv e-prints*, 2007.
- [46] F. Nobile, R. Tempone, and C.G. Webster. An anisotropic sparse grid

- stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.*, 46(5):2411–2442, 2008.
- [47] F. Nobile, R. Tempone, and C.G. Webster. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.*, 46(5):2309–2345, 2008.
- [48] A. Nouy. A generalized spectral decomposition technique to solve a class of linear stochastic partial differential equations. *Computer Methods in Applied Mechanics and Engineering*, 196(37-40):4521–4537, 2007.
- [49] A. Nouy. Generalized spectral decomposition method for solving stochastic finite element equations: Invariant subspace problem and dedicated algorithms. *Computer Methods in Applied Mechanics and Engineering*, 197:4718–4736, 2008.
- [50] M.R. Osborne, B. Presnell, and B. Turlach. A new approach to variable selection in least squares problems. *IMA J. Numer. Anal.*, 20:389–403, 2000.
- [51] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad. Orthogonal Matching Pursuit: Recursive function approximation with applications to wavelet decomposition. In *Proceedings of the 27th Annual Asilomar Conference on Signals, Systems, and Computers*, pages 40–44, 1993.
- [52] C. Schwab and R.A. Todor. Karhunen-Loève approximation of random fields by generalized fast multipole methods. *J. Comput. Phys.*, 217(1):100–122, 2006.
- [53] M.A. Tatang. *Direct incorporation of uncertainty in chemical and environmental engineering systems*. PhD thesis, Ph.D. Thesis, Massachusetts Institute of Technology, 1995.
- [54] R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.
- [55] R. A. Todor and C. Schwab. Convergence rates for sparse chaos approximations of elliptic problems with stochastic coefficients. *IMA Journal of Numerical Analysis*, 27(2):232–261, 2007.
- [56] J.A. Tropp and S.J. Wright. Computational methods for sparse solution of linear inverse problems. *Proceedings of the IEEE*, 2010. in press.
- [57] 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.
- [58] X. Wan and G. Karniadakis. An adaptive multi-element generalized polynomial chaos method for stochastic differential equations. *J. Comp. Phys.*, 209:617–642, 2005.
- [59] R. Ward. Compressed sensing with cross validation. *IEEE Trans. Inf. Theor.*, 55(12):5773–5782, 2009.
- [60] D. Xiu and J.S. Hesthaven. High-order collocation methods for differential equations with random inputs. *SIAM J. Sci. Comput.*, 27(3):1118–1139, 2005.
- [61] D. Xiu and G.M. Karniadakis. The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computin*,

24(2):619–644, 2002.

- [62] W. Xu, M.A. Khajehnejad, A.S. Avestimehr, and B. Hassibi. Breaking through the thresholds: an analysis for iterative reweighted ℓ_1 minimization via the Grassmann Angle Framework. *ArXiv e-prints*, 2009. Available from <http://arxiv.org/abs/0904.0994>.
- [63] J. Yang and Y. Zhang. Alternating direction algorithms for ℓ_1 -problems in compressive sensing. *ArXiv e-prints*, 2009. Available from <http://arxiv.org/abs/0912.1185>.