Coherence Motivated Sampling and Convergence Analysis of Least Squares Polynomial Chaos Regression

Jerrad Hampton, Alireza Doostan*

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

Abstract

Independent sampling of orthogonal polynomial bases via Monte Carlo is of interest for uncertainty quantification of models, using Polynomial Chaos (PC) expansions. It is known that bounding the spectral radius of a random matrix consisting of PC samples, yields a bound on the number of samples necessary to identify coefficients in the PC expansion via solution to a least squares regression problem. We present a related analysis which guarantees a mean square convergence using a *coherence parameter* of the sampled PC basis that may be both analytically bounded and computationally estimated. Utilizing asymptotic results for orthogonal polynomials, we bound the coherence parameter for polynomials of Hermite and Legendre type under each respective natural sampling distribution. In both polynomial bases we identify an importance sampling distribution which yields a bound with weaker dependence on the order of the PC basis. For more general orthonormal bases, we propose the *coherence-optimal* sampling: a Markov Chain Monte Carlo sampling, which directly uses the basis functions under consideration to achieve a statistical optimality among all such sampling schemes with identical support, and which guarantees recovery with a number of samples that is, up to a logarithmic factor, linear in the number of basis functions considered. We demonstrate these different sampling strategies numerically in both high-order and high-dimensional manufactured PC expansions. In addition, the quality of each sampling method is compared in the identification of solutions to two differential equations, one with a high-dimensional

*Corresponding Author: Alireza Doostan

Email address: `alireza.doostan@colorado.edu` (Alireza Doostan)

random input and the other with a high-order PC expansion. In all observed cases the coherence-optimal sampling leads to similar or considerably improved accuracy over the other considered sampling distributions.

Keywords: Polynomial Chaos, Least squares regression, Markov Chain Monte Carlo, Hermite Polynomials, Legendre Polynomials, Uncertainty Quantification

1. Introduction

A precise approach to analyzing complex engineering systems requires understanding how various Quantities of Interest (QoI's) depend upon system inputs that are often uncertain. An incomplete understanding may lead to poor design decisions based upon unfounded sureness or doubt in the QoI's. Uncertainty Quantification (UQ) is an emerging field that aims at addressing these issues in an integral and rigorous manner: it intends for a meaningful characterization of uncertainties from the available information and efficient propagation of these uncertainties for a quantitative validation of model predictions, [1, 2, 3, 4].

Probability is a natural framework for modeling uncertainty, wherein we assume uncertain inputs are represented by a d -dimensional random vector $\Xi := (\Xi_1, \dots, \Xi_d)$ with some joint probability density function $f(\xi)$. Additionally, we assume the coordinates of Ξ are independent, that is $f(\xi)$ is a product of its marginal distributions for each coordinate. In this manner we model the scalar QoI, e.g., a functional of the solution to a physical system and here denoted by $u(\Xi)$, as an unknown function of the inputs, which we seek to approximate. In this work we approximate $u(\Xi)$, assumed to have a finite variance, using a Fourier-type expansion in multivariate orthogonal polynomials, each of which we denote by $\psi_k(\Xi)$, referred to as Polynomial Chaos (PC) expansion [1, 5],

$$u(\Xi) = \sum_{k=0}^{\infty} c_k \psi_k(\Xi). \quad (1)$$

For computation, we allow an arbitrary finite number of input dimensions d but assume u can be accurately approximated in a truncated set of basis of total order less than or equal to p . To briefly explain this total order, let $\mathbf{k} = (k_1, \dots, k_d)$ be a $d \times 1$ multi-index such that $k_i \in \mathbb{N} \cup \{0\}$ represents

the order of the polynomial $\psi_{k_i}(\Xi_i)$ orthonormal with respect to the probability density of Ξ_i . For instance, when Ξ_i follows a uniform or Gaussian distribution, $\psi_{k_i}(\Xi_i)$ are (normalized) Legendre or Hermite polynomials, respectively, [5, 6]. The d -dimensional polynomials $\psi_{\mathbf{k}}(\Xi)$ are constructed by the tensorization of $\psi_{k_i}(\Xi_i)$,

$$\psi_{\mathbf{k}}(\Xi) = \prod_{i=1}^d \psi_{k_i}(\Xi_i). \quad (2)$$

The total order of p implies that we consider all polynomials satisfying

$$\sum_{i=1}^d k_i \leq p \quad k_i \in \mathbb{N} \cup \{0\} \quad \forall i. \quad (3)$$

A direct combinatorial count implies that a d -dimensional approximation of total order p has $P = \binom{p+d}{d}$ basis polynomials. This total order basis facilitates a polynomial approximation to u of the form

$$u(\Xi) \approx \sum_{k=1}^P c_k \psi_k(\Xi), \quad (4)$$

which favors lower order polynomials. The error introduced by this truncation is referred to as *truncation error*, and converges to zero – in the mean squares sense – when $c_k = \mathbb{E}[u(\Xi)\psi_k(\Xi)]$ and p (hence P) $\rightarrow \infty$. Here, $\mathbb{E}$ denotes the mathematical expectation operator.

To identify the PC coefficients $\mathbf{c} = (c_1, \dots, c_P)^T$ in (4), we consider non-intrusive, i.e., sampling-based, methods where we do not require changes to deterministic solvers for u as we generate realizations of Ξ to identify $u(\Xi)$. We denote these realizations $\xi^{(i)}$ and $u(\xi^{(i)})$, respectively. We let $i = 1 : N$ so that N is the number of independent samples considered, and define

$$\begin{aligned} \mathbf{u} &:= (u(\xi^{(1)}), \dots, u(\xi^{(N)}))^T; \\ \Psi(i, j) &:= \psi_j(\xi^{(i)}), \end{aligned} \quad (5)$$

where we refer to Ψ as the *measurement matrix*. These definitions imply the matrix relation $\Psi \mathbf{c} \approx \mathbf{u}$. We also introduce a diagonal positive-definite matrix $\mathbf{W}$ such that $\mathbf{W}(i, i)$ is a function of $\xi^{(i)}$ which depends on our sampling strategy and is described in Sections 2 and 3. To approximate $\mathbf{c}$ we consider

a least squares approximation. Specifically, our theoretical results focus on the solution to the least squares optimization problem,

$$\arg \min_{\mathbf{c}} \|\mathbf{W}\mathbf{u} - \mathbf{W}\Psi\mathbf{c}\|_2, \quad (6)$$

whose solution $\mathbf{c}$ may be computed from the normal equation $(\mathbf{W}\Psi)^T(\mathbf{W}\Psi)\mathbf{c} = (\mathbf{W}\Psi)^T\mathbf{W}\mathbf{u}$, when the matrix $\mathbf{W}\Psi$ is full rank. The study of this problem is useful in its own right, as well as in consideration with the Probabilistic Collocation Method [7, 8, 9, 10], where points are often chosen with (6) in mind. For most of this work, we assume the truncation error model

$$u(\Xi) = \sum_{k=1}^P c_k \psi_k(\Xi) + \epsilon(\Xi). \quad (7)$$

We often do not express ϵ 's dependence on Ξ , which should not be confusing when interpreted in context. Additionally, in Section 2.4 we provide theoretical results pertaining to a more complete error model, which is capable of handling interpretations of model or measurement errors.

The least squares regression problem (6) is among major techniques used to generate PC expansions for uncertainty quantification of engineering systems, see, e.g., [11, 12, 2, 13, 14, 15]. There has been a growing recent interest in understanding the conditions under which problem (6) leads to accurate and stable PC approximation, primarily when the PC basis is of Legendre type and the QoI is sampled according to uniform measure. In particular, [16] presents a bound on the approximation error (in the mean squares sense) for the case of univariate Legendre expansions, provided that the number of samples grows quadratically in the size of the basis. The numerical tests in [17] confirm the need for a quadratic growth of sample size, but reveal that a linear growth may be sufficient for accurate Legendre expansions in multiple dimensions when samples are drawn according to the uniform measure. Bounds on the expected mean squares error were derived in [18, 19] for bounded orthogonal polynomials when sampled according to their natural orthogonality measure. The work in [20] proposes a deterministic sampling strategy that leads to stable approximations in the Chebyshev basis with a number of samples that scales quadratically in the dimension of the polynomial space. The sampling is also extended to other polynomial bases via a weighted least squares formulation. Last but not least, we mention the recent work in [21] that demonstrates empirically the instability of high order

Hermite polynomial expansions and proposes an alternative, stable approximation in the basis of Hermite functions. The QoI (in the univariate case) is evaluated at sample points obtained from a nonlinear mapping of random samples uniformly distributed over $[-1,1]$ to the real line. Under this sampling – which involves parameters to tune – it is shown that the Hermite function expansion is stable when the number of samples that depends linearly on the number of basis functions.

1.1. Contributions of This Work

This work is concerned with the convergence analysis of least squares PC regressions with a particular focus on the role of sampling measure according to which the QoI is evaluated. The novelty of this work arises from utilizing an analysis of Legendre and Hermite polynomials together with the main parameter from [18] to identify sampling methods for stable and accurate solution recovery via (6), as well as direct practical bounds on the number of samples needed to guarantee recovery. As a result, we present concrete theoretical and practical justifications for when a linear (or quadratic) bound on the number of samples is applicable, as well as a sampling that universally guarantees a linear bound, up to a logarithmic factor. In particular, to our best knowledge, the latter sampling results are the first of their type. Additionally, these analyses highlight the role of the order p and dimension d of the PC expansion in the choice of sampling strategy and quality of solution recovery.

In details, we present results demonstrating the stability of least squares solution for *low* order, but possibly *high* dimensional, Legendre and Hermite expansions when sampled according to their standard orthogonality measure, i.e., uniform and standard Gaussian, respectively. Such a stability is guaranteed for a number of samples that scales linearly on the size of the basis set.

Additionally, we provide a contribution of particular practical interest in that we also consider sampling *low* dimensional, but possibly *high* order, Hermite polynomials uniformly over a d -dimensional ball – with a radius depending on the order of approximation – instead of sampling from the standard Gaussian measure. This sampling of Hermite polynomials is analogous to Hermite function expansion, as in [21], when the QoI is appropriately weighted. The uniform sampling of the present work is substantially different from the sampling strategy proposed in [21]. We provide analytic and numerical results justifying the use of this *importance sampling* distribution

for the recovery of Hermite PC expansions. In a similar context, sampling of Legendre polynomials according to Chebyshev measure will be discussed.

Finally, we analytically identify an importance sampling distribution with a statistical *optimality*, in terms of the *spectral radius* of the realized matrix Ψ as a key recovery parameter of the method, and identify a Monte Carlo Markov Chain sampler for which we provide associated numerical results. This approach, here referred to as *coherence-optimal* sampling, provides a general sampling scheme for the reconstruction of expansions in general bases. This sampling scheme allows us to guarantee recovery with a number of samples that scales linearly in the number of basis functions, up to a logarithmic factor, and may be extended to PC types other than Legendre and Hermite. We highlight that such a linear scaling of the number of samples makes the least squares regression approach a competitive framework for UQ of systems where the QoI is not cheap to evaluate.

The presentation in this work has Section 2 providing necessary background information and motivating our approach, while Section 3 describes our sampling methods and provides key theoretical results. Section 4 demonstrates the performance of the sampling methods, and Section 5 presents the proofs to the Theorems from Section 3.

2. Review, Definitions and Recovery Results

To place the results presented in Section 2.3 in context, we provide a summary of recent results on least squares PC regression, while introducing two coherence definitions that are fundamental to our theoretical results, and in constructing our sampling methods.

2.1. Sampling Definitions

We consider as a basis the set of orthonormal polynomials, $\{\psi_k(\xi)\}_{k=1}^P$, and define $B(\xi)$ to be

$$B(\xi) := \sqrt{\sum_{k=1}^P |\psi_k(\xi)|^2}. \quad (8)$$

That is, $B^2(\xi)$ represents a uniformly least upper bound on the sum of the squares of the basis polynomials of interest, and is motivated by work in [18], which uses this function to generate a condition-type number useful for guaranteeing recovery. A similar approach is taken in [22, 23] for the case of

solutions recovered via ℓ_1 -minimization. A bound on $B(\boldsymbol{\xi})$ may be attained from

$$B^2(\boldsymbol{\xi}) \leq P \sup_{k=1:P} |\psi_k(\boldsymbol{\xi})|^2, \quad (9)$$

where $\sup_{k=1:P} |\psi_k(\boldsymbol{\xi})|^2$ may be bounded for several orthogonal polynomials [24, 25, 26, 27, 28, 29, 23]. In addition, to ease analysis of $B(\boldsymbol{\xi})$, formed of piecewise continuous polynomials, we consider computationally convenient functions $G(\boldsymbol{\xi})$ satisfying

$$G(\boldsymbol{\xi}) \geq B(\boldsymbol{\xi}) \quad \forall \boldsymbol{\xi} \in \Omega, \quad (10)$$

which for finite P represents an $L_2(\Omega, f)$ -integrable upper bound on $B(\boldsymbol{\xi})$ over all $\boldsymbol{\xi} \in \Omega$, where Ω denotes the set of possible realizations of $\boldsymbol{\Xi}$.

In this case, we have that $\psi_k(\boldsymbol{\xi})/G(\boldsymbol{\xi}) \leq 1$. It follows that for any set $\mathcal{S} \subseteq \Omega$,

$$f_{\mathbf{Y}}(\boldsymbol{\xi}) := c^2 f(\boldsymbol{\xi}) G^2(\boldsymbol{\xi}), \quad (11)$$

with

$$c = \left(\int_{\mathcal{S}} f(\boldsymbol{\xi}) G^2(\boldsymbol{\xi}) d\boldsymbol{\xi} \right)^{-1/2}, \quad (12)$$

is a probability distribution supported on $\mathcal{S}$, which we consider as the distribution for $\mathbf{Y}$, a copy of $\boldsymbol{\Xi}$ [23]. Let $\delta_{i,j}$ denote the Kronecker delta such that $\delta_{i,j} = 1$ if $i = j$ and 0 if $i \neq j$. Note that for $i, j = 1 : P$,

$$\left| \int_{\mathcal{S}} \frac{\psi_i(\boldsymbol{\xi})}{cG(\boldsymbol{\xi})} \frac{\psi_j(\boldsymbol{\xi})}{cG(\boldsymbol{\xi})} c^2 f(\boldsymbol{\xi}) G^2(\boldsymbol{\xi}) d\boldsymbol{\xi} - \delta_{i,j} \right| \leq \epsilon_{i,j}, \quad (13)$$

and we may select $\mathcal{S}$ such that $\epsilon_{i,j}$ may be made as small as needed, e.g., if we take $\mathcal{S} = \Omega$, then $\epsilon_{i,j} = 0$ due to the orthonormality of $\{\psi_k(\boldsymbol{\Xi})\}_{k=1}^P$. Notice that $\{\psi_k(\mathbf{Y})\}_{k=1}^P$ is no longer (approximately) orthonormal under $f_{\mathbf{Y}}$. Following (13), we therefore consider the weight function

$$w(\mathbf{Y}) := \frac{1}{cG(\mathbf{Y})}, \quad (14)$$

so that $\{w(\mathbf{Y})\psi_k(\mathbf{Y})\}_{k=1}^P$ are (approximately) orthonormal random variables. This function defines the diagonal positive-definite matrix $\mathbf{W}$ from (6) as

$$\mathbf{W}(i, i) = w(\boldsymbol{\xi}^{(i)}), \quad (15)$$

where $\boldsymbol{\xi}^{(i)}$ is the i th realization of $\mathbf{Y}$. For the interest of notation, we refer to all realized random vectors by $\boldsymbol{\xi}$ regardless of the sampling distribution for $\boldsymbol{\xi}$, noting that the weight function w depends on that distribution. Additionally, we note that for simulation, we are typically not interested in the normalizing constant c associated with the sampling distribution.

2.2. Coherence Parameter Definition

Consider realizations of $w(\mathbf{Y})\psi_k(\mathbf{Y})$ for $k = 1 : P$. We investigate the *coherence* parameter

$$\mu_\Omega(\mathbf{Y}) := \sup_{\boldsymbol{\xi} \in \Omega} \sum_{k=1}^P |w(\boldsymbol{\xi})\psi_k(\boldsymbol{\xi})|^2, \quad (16)$$

which, following [18], yields a bound on the realized spectral radius of $\mathbf{W}\Psi$ and useful results pertaining to the solution of (6). We often suppress the dependence of μ_Ω on the sampling scheme, as there should be no confusion that it depends on how the random variables are sampled. From (14) we are motivated to consider $G(\boldsymbol{\xi})$ taken to be $B(\boldsymbol{\xi})$ as defined in (8), and we shall show in Theorem 3.3 that this sampling method leads to $\mu_\Omega = P$ and a linear dependence on the number of samples needed as compared to the number of basis functions considered up to a logarithmic factor in P . Fortunately, asymptotic (in p) results on the orthogonal polynomials considered here give choices for $G(\boldsymbol{\xi})$ that, while generally lead to sub-optimal μ_Ω , correspond to sampling from known measures $f_{\mathbf{Y}}$, as discussed in Section 3.2.

We utilize the definition in (16) when analyzing Legendre polynomials which are bounded on the domain $[-1, 1]^d$. However, we note that (16) is not useful when $\sup_{k=1:P, \boldsymbol{\xi} \in \Omega} |w(\boldsymbol{\xi})\psi_k(\boldsymbol{\xi})|^2$ is infinite, such as when $\psi_k(\boldsymbol{\xi})$ are Hermite polynomials and $w(\boldsymbol{\xi}) = 1$. In truth we may relax the conditions to a supremum not over Ω , but a sufficiently large subset of Ω . If N is the number of samples of $\mathbf{Y}$ which we will take, motivated by the truncation in [22, 23] for recovery via ℓ_1 -minimization, we consider a truncation of Ω to an appropriately measurable subset $\mathcal{S}$ of Ω . This leads to a definition of

coherence given by

$$\mu_{\mathcal{S}}(\mathbf{Y}) := \begin{cases} \underset{\mathcal{S}}{\text{minimize}} & \sup_{\boldsymbol{\xi} \in \mathcal{S}} \sum_{k=1}^P |w(\boldsymbol{\xi})\psi_k(\boldsymbol{\xi})|^2, \\ \text{subject to} & \mathbb{P}(\mathcal{S}^c) \leq \frac{1}{NP}, \\ & \text{and} \quad \sum_{k=1}^P \mathbb{E}(|w(\mathbf{Y})\psi_k(\mathbf{Y})|^2 \mathbf{1}_{\mathcal{S}^c}) \leq \frac{1}{20} P^{-1/2}, \end{cases} \quad (17)$$

where the superscript c denotes a set complement and $\mathbf{1}$ is an indicator function. As before, the dependence on the sampling scheme is often suppressed. If $\mathcal{S}$ is a finite set, then $\mu_{\mathcal{S}}$ is finite for polynomial basis functions, as we consider here. For the purpose of truncating to $\mathcal{S}$ we employ the heuristic of selecting $\mathcal{S}$ to encompass the largest values of $f(\boldsymbol{\xi})$ until $\mathcal{S}$ is large enough to satisfy the conditions in (17). The justification for this is that in unbounded domains, e.g., for Hermite polynomials, regions of small $f(\boldsymbol{\xi})$ typically correspond to larger $B(\boldsymbol{\xi})$ as p grows [25, 27, 26].

Remark. We note that the conditions in (17) can be tweaked, with tradeoffs concerning the probability of failure of (6) in finding a solution $\mathbf{c}$ and a bias in spectral norm of $(\mathbf{W}\Psi)^T(\mathbf{W}\Psi)$ as discussed in Section 5. Also, for the sampling described in Section 3.3, we can forego this truncation. We also consider the same truncation to $\mathcal{S}$ of (17) in regards to the coherence parameter $\mu_{\infty}(\mathbf{Y})$ defined by

$$\mu_{\infty}(\mathbf{Y}) := \sup_{\boldsymbol{\xi} \in \mathcal{S}, k=1:P} |w(\boldsymbol{\xi})\psi_k(\boldsymbol{\xi})|^2. \quad (18)$$

This parameter is often easier to directly bound than the sum associated with $\mu_{\mathcal{S}}$, or μ_{Ω} . Notice that μ_{∞} may be used to bound $\mu_{\mathcal{S}}$ via

$$\mu_{\mathcal{S}}(\mathbf{Y}) \leq P\mu_{\infty}(\mathbf{Y}). \quad (19)$$

This bound may be sharp, for example when $|w(\boldsymbol{\xi})\psi_k(\boldsymbol{\xi})|^2 = 1$ uniformly in $\boldsymbol{\xi}$ and k , as occurs when using basis functions restricted to have complex modulus one for $\boldsymbol{\xi} \in \mathcal{S}$. Additionally, one has that $\mu_{\mathcal{S}} \geq \mu_{\infty}$, which may also be sharp if, for instance, the basis functions have non-overlapping support. However, these cases do not apply to the polynomial bases considered here. In fact, as has been investigated in [23], and will be reproduced in Section 3 in regards to polynomial bases, with an appropriate sampling distribution μ_{∞} is most heavily influenced by p , the order of PC approximation, or d , the dimension of approximation.

As seen in the next section, following [18], solution recovery from (6) can be guaranteed with a number of samples that is proportional to μ_S up to a logarithmic factor in P . This is bounded by a number of samples that is proportional to $\mu_\infty P \log(P)$. The sampling distribution f_Y will imply the coherence μ_∞ and its dependence on d and p . In particular, in the case of Legendre polynomials, μ_∞ grows exponentially in p under standard sampling but is bounded in d . Similarly, μ_∞ grows exponentially in d under Chebyshev sampling, but is bounded in p . See for example the results of [23], whose relevant theorems for bounds on μ_∞ are presented in Section 3. We note that this analysis can be directly used to explain observations regarding solution recovery from (6) using a number of samples that depends either linearly or quadratically on the number of basis functions P , see, e.g., [19, 16, 15, 21].

2.3. Convergence Theorems

The following theorems use the coherence parameter in either (16) or (17) to bound the number of samples necessary to recover the PC coefficients from (6) with high probability. Our first theorem concerning solution recovery is comparable to Theorem 2 in [18], and demonstrates a quality mean squared convergence with high probability.

Theorem 2.1. *Let*

$$\hat{u}(\Xi) = \sum_{k=1}^P \hat{c}_k \psi_k(\Xi),$$

where $\hat{\mathbf{c}} = (\hat{c}_1, \dots, \hat{c}_P)^T$ is the solution to (6). For $\mathcal{E}$ a sampling event (defined in Section 5) that occurs with probability

$$\mathbb{P}(\mathcal{E}) \geq 1 - 1/P - 2P \exp(-0.1N\mu_S^{-1}),$$

it follows that

$$\mathbb{E} \left(\|u - \hat{u}\|_{L_2(\Omega, f)}^2; \mathcal{E} \right) \leq \mathbb{E}(\epsilon^2) \left(1 + \frac{4\mu_S}{N} \right),$$

where

$$\mathbb{E}(X; \mathcal{E}) = \int_{\mathcal{E}} X(\xi) f(\xi) d\xi = \mathbb{E}(X|\mathcal{E}) \mathbb{P}(\mathcal{E}) \quad (20)$$

denotes the expectation restricted to the event (also known as restricted expectation), and is closely related to conditional expectation. Similarly, with $\mathcal{E}$ occurring with probability

$$\mathbb{P}(\mathcal{E}) \geq 1 - 2P \exp(-0.1N\mu_\Omega^{-1}),$$

it follows that

$$\mathbb{E}(\|u - \hat{u}\|_{L_2(\Omega, f)}^2; \mathcal{E}) \leq \mathbb{E}(\epsilon^2) \left(1 + \frac{4\mu_\Omega}{N}\right).$$

Remark. We note that expectations in this theorem – as well as in the subsequent theorems – are with respect to the original measure $f(\boldsymbol{\xi})$ of $\boldsymbol{\xi}$, even when we sample according to $f_{\mathbf{Y}}$, see the discussion in Section 5.1.

The Theorem 2.1 and its dependence on $\mu_{\mathcal{S}}$ motivate us to sample such that $\mu_{\mathcal{S}}$ is small. Indeed, if $\mu_{\mathcal{S}}$ scales linearly in P , as for example under the sampling of Section 3.3 or when μ_∞ is bounded by a constant, then we can take N to scale nearly linearly with P in a way that depends on the desired accuracy. Recalling that ϵ is the residual arising from an orthogonal projection of u onto the polynomial basis, and thus the smallest possible $L_2(\Omega, f)$ error attainable by an approximation, it follows that $\|u - \hat{u}\|_{L_2(\Omega, f)}^2 - \mathbb{E}(\epsilon^2) \geq 0$. Applying the Markov inequality, gives a corollary that is also useful in bounding a failure probability for a particular least squares computation.

Corollary 1. For any $\tau > 0$,

$$\mathbb{P}(\|u - \hat{u}\|_{L_2(\Omega, f)}^2 \geq \mathbb{E}(\epsilon^2) + \tau) \leq \frac{1}{P} + 2P \exp\left(-\frac{0.1N}{\mu_{\mathcal{S}}}\right) + \frac{4\mu_{\mathcal{S}}}{N} \cdot \frac{\mathbb{E}(\epsilon^2)}{\tau}.$$

Further, when $\mu_{\mathcal{S}}$ is replaced by μ_Ω , this bound holds without the $1/P$ term.

Additionally, the following theorem is useful for bounding needed sample sizes follows directly from Theorem 2.1,

Theorem 2.2. For any $\tau > 0$, let

$$\tau N_\tau \geq 4\mathbb{E}(\epsilon^2)\mu_{\mathcal{S}},$$

using the same definitions as in Theorem 2.1. It follows that

$$\mathbb{E}(\|u - \hat{u}\|_{L_2(\Omega, f)}^2; \mathcal{E}) \leq \mathbb{E}(\epsilon^2) + \tau, \tag{21}$$

where $\mathcal{E}$ occurs with probability

$$\mathbb{P}(\mathcal{E}) \geq 1 - 1/P - 2P \exp(-0.1N\mu_S^{-1}).$$

Similarly, if

$$\tau N_\tau \geq 4\mathbb{E}(\epsilon^2)\mu_\Omega,$$

then it follows that (21) holds with probability

$$\mathbb{P}(\mathcal{E}) \geq 1 - 2P \exp(-0.1N\mu_\Omega^{-1}).$$

Stated another way, Theorem 2.2 allows us to identify the number of samples N_τ necessary to guarantee with a bounded probability that the L_2 error in the recovery of u is within τ of optimal. We see that if μ_S is such that $\mu_S \leq \nu P$ for some constant $1 \leq \nu \leq \mu_\infty$ depending on the measure of $\mathbf{Y}$, then N_τ scales linearly in P , the number of basis functions considered. This is the case for example, when we sample as described in Section 3.3, or when μ_∞ is bounded by a constant. If we make this assumption that $\mu_S \leq \nu P$, then from the result of Theorem 2.2, the failure probability $1 - \varrho$ is bounded by

$$1 - \varrho \leq \exp\left(\log(2P) - \frac{0.1N}{\nu P}\right).$$

An algebraic rearrangement of terms gives that

$$N \geq 10\nu P \log\left(\frac{2P}{1 - \varrho}\right).$$

This provides an explicit bound on the number of samples needed to guarantee an accurate recovery within a certain failure probability. We summarize this implication and that of error control suggested by Theorem 2.2 in the following corollary.

Corollary 2. *Let μ_Ω be bounded by νP for some $0 < \nu \leq \mu_\infty$ depending on the measure of $\mathbf{Y}$. Then, to guarantee recovery to within τ of minimum mean-squared-error with probability at least ϱ , one can generate a number of samples $N_{\tau,\varrho}$ satisfying*

$$N_{\tau,\varrho} \geq \nu \cdot \max\left\{\frac{4\mathbb{E}(\epsilon^2)}{\tau} \cdot P, 10P \cdot \log\left(\frac{2P}{1 - \varrho}\right)\right\},$$

which scales linearly in P up to a logarithmic factor. If μ_S is used and ϱ is chosen so that $\varrho < 1 - 1/P$, then

$$N_{\tau, \varrho} \geq \nu \cdot \max \left\{ \frac{4\mathbb{E}(\epsilon^2)}{\tau} \cdot P, 10P \cdot \log \left(\frac{2P}{1 - 1/P - \varrho} \right) \right\}.$$

We note that the separation of the two sample size requirements within the max operator of Corollary 2 is fruitful in that each derivation is sharp with regards to the order for the number of samples required, although the multiplying constants are not necessarily optimal.

The left term in the maximum is related to recovery when a quality sampling has been achieved. For general ℓ_2 recovery with P unknowns, because of the the existence of a non-trivial null-space with as many as $P - 1$ samples, a linear growth in the number of samples is the minimum order that is anticipated for this term.

The right term in the maximum is related to identifying that a quality sampling has occurred with probability ϱ . The logarithmic term for this bound is also tight in certain contexts, and is related to the *coupon collector* problem (that is anticipating how many draws from a uniform distribution on $\{1, \dots, P\}$ are needed to have a draw from each of $\{1, \dots, P\}$). We refer the interested reader to Remark 5.6 of [30].

This decoupling of constants is useful for bounding the accuracy of a computed solution while guaranteeing a certain probability of successful recovery, or similarly, identifying a probability of successful recovery while guaranteeing a certain accuracy in the computed solution.

2.4. A More Robust Noise Model

In this section we consider u as a function of (Ξ, χ) , where χ are unobserved random variables that may have parameters which are functions of realized ξ , and are useful to account for model and measurement errors. As we do not observe χ , we do not attempt to compute an estimate $\hat{u}$ that depends on χ and instead attempt to identify the conditional random variable,

$$u(\Xi) := \mathbb{E} \left(u(\Xi, \chi) | \Xi \right), \quad (22)$$

noting that there should be no confusion in the notation when function arguments are present. We then consider the error model,

$$u(\Xi, \chi) := u(\Xi) + \epsilon_M(\chi|\Xi); \quad (23)$$

$$u(\Xi) := \sum_{k=1}^P c_k \psi_k(\Xi) + \epsilon_T(\Xi), \quad (24)$$

where ϵ_T is still interpreted as a truncation error that depends only on Ξ . This model gives us the following theorem, which is comparable to but significantly different than Theorem 3 in [18].

Theorem 2.3. *Using the same definitions as in Theorem 2.1, it follows that*

$$\mathbb{E} (\|u - \hat{u}\|_{L_2(\Omega, f)}^2; \mathcal{E}) \leq \mathbb{E}(\epsilon_T^2) + 8\mu_S \frac{\mathbb{E}(\epsilon_T^2) + \mathbb{E}(\epsilon_M^2)}{N},$$

where

$$\mathbb{P}(\mathcal{E}) \geq 1 - 1/P - 2P \exp(-0.1N\mu_S^{-1}).$$

Similarly,

$$\mathbb{E} (\|u - \hat{u}\|_{L_2(\Omega, f)}^2; \mathcal{E}) \leq \mathbb{E}(\epsilon_T^2) + 8\mu_\Omega \frac{\mathbb{E}(\epsilon_T^2) + \mathbb{E}(\epsilon_M^2)}{N},$$

where

$$\mathbb{P}(\mathcal{E}) \geq 1 - 2P \exp(-0.1N\mu_\Omega^{-1}).$$

This theorem demonstrates that the recovery in this case is quite similar to that of Theorem 2.1. Applying the Markov inequality similarly gives a bound on failure probability in this case.

Corollary 3. *For any $\tau > 0$,*

$$\mathbb{P} (\|u - \hat{u}\|_{L_2(\Omega, f)}^2 \geq \mathbb{E}(\epsilon^2) + \tau) \leq \frac{1}{P} + 2P \exp \left(-\frac{0.1N}{\mu_S} \right) + \frac{8\mu_S}{N} \cdot \frac{\mathbb{E}(\epsilon^2) + \mathbb{E}(\epsilon_M^2)}{\tau}.$$

Further, when μ_S is replaced by μ_Ω this bound holds without the $1/P$ term.

Additionally, the number of samples required to guarantee an accurate recovery in expectation is also similar to that of Theorem 2.2, stated in the following theorem, which follows directly from Theorem 2.3.

Theorem 2.4. For any $\tau > 0$, let

$$\tau N_\tau \geq 8\mu_S (\mathbb{E}(\epsilon_T^2) + \mathbb{E}(\epsilon_M^2)) ,$$

using the same definitions as in Theorem 2.1. It follows that

$$\mathbb{E} (\|u - \hat{u}\|_{L_2(\Omega, f)}^2; \mathcal{E}) \leq \mathbb{E}(\epsilon_T^2) + \tau, \quad (25)$$

where $\mathcal{E}$ occurs with probability

$$\mathbb{P}(\mathcal{E}) \geq 1 - 1/P - 2P \exp(-0.1N\mu_S^{-1}).$$

Similarly, if

$$\tau N_\tau \geq 8\mu_\Omega (\mathbb{E}(\epsilon_T^2) + \mathbb{E}(\epsilon_M^2)) ,$$

then it follows that (25) holds with probability

$$\mathbb{P}(\mathcal{E}) \geq 1 - 2P \exp(-0.1N\mu_\Omega^{-1}).$$

We can make a similar explicit bound on the number of samples for this noise model as well.

Corollary 4. Let μ_Ω be bounded by νP for some $0 \leq \nu \leq \mu_\infty$ depending on the measure of $\mathbf{Y}$. Then to insure recovery to within τ of minimum mean-squared-error with probability at least ϱ , one can generate a number of samples $N_{\tau, \varrho}$ satisfying

$$N_{\tau, \varrho} \geq \nu \cdot \max \left\{ \frac{8 (\mathbb{E}(\epsilon_T^2) + \mathbb{E}(\epsilon_M^2))}{\tau} \cdot P, 10P \cdot \log \left(\frac{2P}{1 - \varrho} \right) \right\} ,$$

which scales linearly in P up to a logarithmic factor. If μ_S is used, and ϱ is chosen so that $\varrho < 1 - 1/P$, then

$$N_{\tau, \varrho} \geq \nu \cdot \max \left\{ \frac{8 (\mathbb{E}(\epsilon_T^2) + \mathbb{E}(\epsilon_M^2))}{\tau} \cdot P, 10P \cdot \log \left(\frac{2P}{1 - 1/P - \varrho} \right) \right\} .$$

Remark. We note that the same considerations as followed Corollary 2 apply here in this slightly more general model, with the only difference being a doubling of the first multiplicative constant.

3. Sampling Methods

Here, we describe the sampling methods that we consider in this work, and present theorems related to recovery when we use them. Our objective is to investigate the implication of certain choices of sampling measure $f_{\mathbf{Y}}$ on coherence, and hence on the parameter ν that controls the bound on sampling rates in Corollaries 2 and 4.

We first consider a sampling according to random variables defined by the orthogonality measure in Section 3.1. Such a sampling, dubbed here *standard sampling*, is commonly used in PC regression, [11, 2, 31, 32]. We also consider sampling from a distribution associated with an asymptotic analysis, although not necessarily corresponding with the asymptotic distribution, of the orthogonal polynomials $\psi_k(\Xi)$ in Section 3.2, and refer to it as *asymptotic sampling*. For the above two sampling methods, we discuss the dependence of μ_∞ , an upper bound on ν , on the dimensions d of inputs and total order p of the PC basis. Finally, in Section 3.3, we introduce the *coherence-optimal sampling* that corresponds to minimizing ν , as opposed to μ_∞ , directly. The empirical results of Section 4 suggest that sampling strategies with (considerably) smaller coherence, lead also to improved solution recoveries, hence signifying the role of the coherence parameter in the actual performance of the least squares PC regression. We note that the proofs for all but the last of the theorems in this section may be found with no changes within [23].

3.1. Standard Sampling

We first consider sampling ξ according to the orthogonality measure $f(\xi)$ of the PC bases, for which we set the weights $w(\xi) = 1$. For the d -dimensional Legendre polynomials the standard method corresponds to sampling from the uniform distribution on $[-1, 1]^d$, while for d -dimensional Hermite polynomials this corresponds to samples from a multi-variate normal distribution such that each of d coordinates is an independent standard normal random variable.

For the standard sampling of orthogonal polynomials of arbitrary, total order p in dimension d , the following results hold for μ_∞ (a bound on the parameter ν in theorems of Section 2).

Theorem 3.1. ([23] Theorems 4.1 and 4.2) *A standard sampling of the d -dimensional Legendre polynomials of total order p gives a coherence of*

$$\mu_\infty(\Xi) \leq \exp(2p). \quad (26)$$

Assume that $d = o(p)$, and let $N = O(P^k)$ for some positive, finite k . For d -dimensional Hermite polynomials of total order $p \geq 1$, the coherence in (18) is bounded by

$$\mu_\infty(\Xi) \leq C_p \cdot \eta_p^p, \quad (27)$$

for constants C_p, η_p depending on p . For $d = o(p)$, and as $p \rightarrow \infty$, we may take C_p and η_p to be larger than but arbitrarily close to 1 and $\exp(2 - \log(2)) \approx 3.6945$, respectively.

As shown in [23], for the case of $p < d$ the bound in (26) may be improved to $\mu_\infty(\Xi) \leq 3^p \approx \exp(1.1p)$. Additionally, for $p > d$, the bound in (26) is loose, but a sharper dimension-dependent bound is given by $(2p/d + 1)^d$.

Remark. When sampling Hermite polynomials, we have the technical requirement that $N = O(P^k)$ for some finite k . We note that this includes the sample sizes suggested by the corollaries in Section 2.

We note that the combination of the bounds $\nu \leq \mu_\infty$ and (27), together with Corollary 2, implies that the number of samples required for recovery from Hermite polynomials may grow exponentially with the total order of approximation.

3.2. Asymptotic Sampling

We next consider taking $G(\xi)$ to be an asymptotic envelope for the polynomials as the order p goes to infinity. For very large p , it then follows that $G(\xi)$ is an approximation to $B(\xi)$.

For the d -dimensional Hermite polynomials, an analysis that utilizes Hermite Functions and presented in [23] leads to a choice of $G(\xi)$ corresponding to sampling uniformly from within the d -dimensional ball of radius $\sqrt{2}\sqrt{2p+1}$. This leads to a weight function given by

$$w(\xi) := \exp(-\|\xi\|_2^2/4). \quad (28)$$

Following [23], the algorithm below may be utilized to sample uniformly from the d -dimensional ball of radius r . First, let $\mathbf{Z} := (Z_1, \dots, Z_d)$ be a vector of d independent normally distributed random variables with zero mean and the same variance. If U is another independent random variable that is uniformly distributed on $[0, 1]$, then

$$\mathbf{Y} := \frac{\mathbf{Z}}{\|\mathbf{Z}\|_2} r U^{1/d}, \quad (29)$$

represents a random sample uniformly distributed within the d -dimensional ball of radius r .

For the d -dimensional Legendre polynomials the asymptotic sampling measure is known as the Chebyshev distribution on $[-1, 1]^d$, [24], that is the distribution in each of the d coordinates is

$$f_Y(\xi) := \frac{1}{\pi\sqrt{1-\xi^2}}, \quad (30)$$

for $\xi \in [-1, 1]$. Each coordinate is easily simulated from $\cos(\pi U)$ where U is uniformly distributed on $[0, 1]$. Additionally, this leads to a weight function given by

$$w(\boldsymbol{\xi}) := \prod_{i=1}^d (1 - \xi_i^2)^{1/4}. \quad (31)$$

In the following theorem we bound the coherence parameter μ_∞ associated with these alternative sampling of Legendre and Hermite polynomials, and demonstrate a weaker asymptotic dependence of μ_∞ on p .

Theorem 3.2. ([23] Theorems 4.3 and 4.4) *For the sampling of d -dimensional Legendre polynomials according to the d -dimensional Chebyshev distribution and weight $\psi_k(\boldsymbol{\xi})$ proportional to $w(\boldsymbol{\xi}) = \prod_{i=1}^d (1 - \xi_i^2)^{1/4}$, regardless of the relationship between d and p , we have that*

$$\mu_\infty \leq 3^d. \quad (32)$$

Assume that $N = O(P^k)$ for some finite $k > 0$, which includes the sample sizes suggested by the Theorems in Section 2. Let $V(r, d) = (r\sqrt{\pi})^d / \Gamma(d/2 + 1)$ denote the volume inside the hypersphere with radius r in dimension d .

For the sampling of Hermite polynomials, sampling uniformly from the d -dimensional ball of radius $\sqrt{2}\sqrt{(2 + \epsilon_p)p + 1}$, and weighting realized $\psi_k(\boldsymbol{\xi}^{(i)})$ on this ball by $w(\boldsymbol{\xi}^{(i)}) = \exp(-\|\boldsymbol{\xi}^{(i)}\|_2^2/4)$, gives

$$\mu_\infty = O(\pi^{-d/2} V(\sqrt{2p}, d)) = O((2p)^{d/2} / \Gamma(d/2 + 1)). \quad (33)$$

We note that for Legendre polynomials sampled by Chebyshev distribution we have a complete independence of the order of approximation, which agrees with previous results in [24]. We also highlight that the results of Theorems 3.1 and 3.2 suggest sampling Hermite and Legendre polynomials from

the standard distribution when $d > p$ and from the asymptotic distribution when $d < p$. A similar observation has been made in [33].

Remark. In the numerical results of this work, for uniform sampling of Hermite polynomials, we set ϵ_p in Theorem 3.2 to be zero, leaving as an open problem the determination of an optimal ϵ_p , and hence sampling radius.

3.3. Coherence-optimal Sampling

We finally consider taking $G(\boldsymbol{\xi}) = B(\boldsymbol{\xi})$ in (11), which implies sampling $\boldsymbol{\xi}$ according to the distribution

$$f_{\mathbf{Y}}(\boldsymbol{\xi}) = c^2 B^2(\boldsymbol{\xi}) f(\boldsymbol{\xi}), \quad (34)$$

where $f(\boldsymbol{\xi})$ is the measure with respect to which $\{\psi_k(\boldsymbol{\xi})\}_{k=1}^P$ are naturally orthogonal. Recalling the definition of $B^2(\boldsymbol{\xi})$ as the sum of the squares of the basis polynomials of interest, and having $\mathcal{S}$ satisfy (17) implies that the normalizing constant c^2 in (34) is bounded via the relation,

$$P - \frac{1}{20\sqrt{P}} \leq \int_{\mathcal{S}} B^2(\boldsymbol{\xi}) f(\boldsymbol{\xi}) d\boldsymbol{\xi} = c^{-2} \leq \int_{\Omega} B^2(\boldsymbol{\xi}) f(\boldsymbol{\xi}) d\boldsymbol{\xi} = P. \quad (35)$$

Corresponding to this sampling, we apply the weight function

$$w(\boldsymbol{\xi}) = \frac{1}{cB(\boldsymbol{\xi})}, \quad (36)$$

or $w(\boldsymbol{\xi}) = 1/B(\boldsymbol{\xi})$ in practice.

3.3.1. Sampling via Markov Chain Monte Carlo

To avoid computing the normalization constant c in (34), we sample Ξ from (34) using a Monte Carlo Markov Chain (MCMC) approach, specifically via the Metropolis-Hastings sampler [34], which requires evaluations of $B(\boldsymbol{\xi})$ from (8). Additionally, this sampling distribution allows the easy evaluation of $w(\boldsymbol{\xi})$ using only the realized samples. Theorem 3.3 justifies the intuition that taking $G(\boldsymbol{\xi})$ associated with sampling to be the envelope function $B(\boldsymbol{\xi})$ leads to the minimal $\nu = 1$.

Theorem 3.3. *Let $\mathcal{S}$ be a set chosen to satisfy the conditions of (17) such that no subset $\mathcal{T}$ of $\mathcal{S}$ with $\mu_{\mathcal{T}} < \mu_{\mathcal{S}}$ satisfies the conditions of (17). Let $B(\boldsymbol{\xi})$ be as in (8). If we sample from the distribution proportional to $f(\boldsymbol{\xi})B^2(\boldsymbol{\xi})$ and weight $\psi_k(\boldsymbol{\xi})$ proportional to $w(\boldsymbol{\xi}) = 1/B(\boldsymbol{\xi})$, then the coherence parameter*

$\mu_{\mathcal{S}}$ achieves a minimum over all sampling schemes of $\psi_k(\boldsymbol{\xi})$, $k = 1 : P$, and distributions supported on $\mathcal{S}$.

Further, for this sampling strategy, $\mu_{\Omega} = P$.

In more detail regarding the implementation, the MCMC sampler requires a candidate distribution and when $p \leq d$, we suggest those obtained from Section 3.1, giving a standard normal sampling for Hermite polynomials, and sampling uniformly for $[-1, 1]^d$ for Legendre polynomials. Similarly, when $p > d$ we suggest those obtained from Section 3.2, giving a uniform sampling on a d -dimensional ball for Hermite polynomials, and d -dimensional Chebyshev sampling for Legendre polynomials. Note that each proposal distribution covers the entire domain $\mathcal{S}$, and if the proposal and target distribution approximately match, then the acceptance rate is high and few burn-in samples are needed to approximately draw from the desired distribution for $\mathbf{Y}$. One caveat which we note is that the proofs of Theorems 2.1 and 2.3 require independent sampling, so that it is proper to restart a chain after each accepted sample, but a more practical method is to discard intermediate samples so that serial dependence is small, [35]. We note that in applications where evaluation of the QoI is computationally expensive, the generation of the samples $\{\boldsymbol{\xi}^{(i)}\}_{i=1}^N$ is not typically a bottleneck, so that the extra cost for quality MCMC sampling is frequently acceptable in practice. Algorithm 1 provides an outline of the steps involved in generating coherence-optimal samples of Ξ according to $f_{\mathbf{Y}}$ in (34).

4. Numerical Examples

We next study the empirical performance of different sampling schemes discussed in Section 3, first exploring an estimate of the coherence parameters in Section 4.1. We then investigate the recovery of randomly generated functions in Section 4.2, and functions given by the solution to an ODE with random coefficient modeling the amount of reaction at a given time in an adsorption model from [36] in Section 4.3. Additionally, we consider a high-dimensional model of a thermally driven cavity flow in Section 4.4.

4.1. Estimates of Coherence

The coherence parameters of Section 2.2 can be estimated using a large sample of realized $\psi_k(\boldsymbol{\xi})$. Doing so leads to the results in Figure 1 for Hermite polynomials, and those in Figure 2 for Legendre polynomials. We consider

Algorithm 1 MCMC sampler for coherence-optimal sampling from the density in (34).

Select a proposal distribution, denoted $g(\boldsymbol{\xi})$, based on distribution of Ξ and whether $p > d$ or $p \leq d$, as described in Section 3.3.1.
 Draw $\boldsymbol{\xi}$ from $g(\boldsymbol{\xi})$.
for A number of burn-in samples **do**
 Draw $\boldsymbol{\xi}'$, independently, from the proposal distribution.
 Set $\rho := \min \left\{ 1, \frac{g(\boldsymbol{\xi})f(\boldsymbol{\xi}')B^2(\boldsymbol{\xi}')}{g(\boldsymbol{\xi}')f(\boldsymbol{\xi})B^2(\boldsymbol{\xi})} \right\}$.
 Draw U , independently, from $[0, 1]$.
 if $U < \rho$ **then**
 Set $\boldsymbol{\xi} = \boldsymbol{\xi}'$.
 end if
end for
 $\boldsymbol{\xi}$ is approximately drawn from (34).

three sampling schemes, the standard scheme where we sample based on the underlying distribution of random variables in question as in Section 3.1; an asymptotically motivated method to insure a coherence with weaker dependence on the order p as in Section 3.2; and a coherence-optimal sampling based on the distribution proportional to the envelope of basis functions as in Section 3.3. We observe that standard sampling has larger coherence at high orders, while asymptotic sampling has larger coherence for high-dimensional problems. Additionally, coherence-optimal sampling performs well in all regimes. These observations are consistent with the theoretical results presented in Section 3.

4.2. Manufactured Functions

In this section, we investigate the reconstruction accuracy of the competing sampling schemes on randomly generated solution vectors, $\mathbf{c}$, such that $\Psi \mathbf{c} = \mathbf{u}$. Here, each coordinate of $\mathbf{c}$ is an independent standard normal random variable, which implies that, unlike for most solutions of practical interest, the coordinates of $\mathbf{c}$ do not exhibit a decay with respect to the order of $\psi_k(\boldsymbol{\xi})$. We measure reconstruction accuracy as a function of the number of independent samples of $\mathbf{Y}$, denoted by N . We declare $\hat{\mathbf{c}}$ to successfully recover $\mathbf{c}$ if $\|\hat{\mathbf{c}} - \mathbf{c}\|_2 / \|\mathbf{c}\|_2 \leq 0.02$, where $\hat{\mathbf{c}}$ is a solution to (6) and in this work is computed using MATLAB's `mldivide` function, which is more ubiquitously known as the backslash (`\`) matrix operation. Each success probability is

Figure 1: Computed coherence for different sampling methods of Hermite polynomials for different d and p .

calculated from an average of a large number of independent realizations of Ψ and $\mathbf{c}$. We consider two cases, one in which there is no noise in the evaluation of $\mathbf{u}$, and one in which there is independent, normally distributed additive noise in the evaluations of $u(\xi^{(i)})$, each having standard deviation $0.03 \cdot |u(\xi^{(i)})|$.

The results in Figures 3 and 4 identify a sharp transition in the ability of the least squares problem (6) to recover $\mathbf{c}$ in the noisy cases. In particular, we highlight the following notable observations: for the high order case $(d, p) = (2, 30)$, the standard sampling fails to recover the solution, while the asymptotic sampling succeeds. For the high-dimensional case $(d, p) = (30, 2)$, the standard sampling is better than the asymptotic sampling, and for the moderate values of $(d, p) = (5, 5)$ the coherence-optimal method outperforms both asymptotic and standard sampling, while the standard sampling of Hermite expansion performs poorly. In all cases, the coherence-optimal sampling

Figure 2: Computed coherence for different sampling methods of Legendre polynomials for different d and p .

leads to recovery that is similar to those of the other two sampling strategies or provides considerable improvements. For the noisy cases, the recovery of the sampling strategies examined here is similar except for the Hermite expansion with $(d, p) = (2, 30)$ where the standard sampling fails to converge within the range of sample sizes considered.

4.3. A Surface Reaction Model with Random Input

The second problem of interest in this work is to quantify the uncertainty in the solution ρ of the non-linear evolution equation

$$\begin{cases} \frac{d\rho}{dt} = \alpha(1 - \rho) - \gamma\rho - \kappa(1 - \rho)^2\rho, \\ \rho(t = 0) = 0.9, \end{cases} \quad (37)$$

which models the surface coverage of certain chemical species, as examined in [36, 37, 23]. We consider uncertainty in the adsorption, α , and desorption, γ , coefficients, and model both as shifted log-normal random variables.

Figure 3: Hermite recovery probability as a function of sample size N : Top graphs are for noisy recovery, while bottom are for noiseless recovery. Horizontal position corresponds to differing (d, p) .

Specifically, we assume

$$\begin{aligned}\alpha &= 0.1 + \exp(0.05 \Xi_1), \\ \gamma &= 0.001 + 0.01 \exp(0.05 \Xi_2),\end{aligned}\tag{38}$$

where Ξ_1, Ξ_2 are independent standard normal random variables; hence, the dimension of our random input is $d = 2$. The reaction rate constant $\kappa = 10$ in (37) is assumed to be deterministic.

Our QoI is $\rho_c := \rho(t = 4, \Xi_1, \Xi_2)$. In Figure 5, we display the realizations of ρ_c at $N = 3000$ samples of (Ξ_1, Ξ_2) generated according to the standard,

Figure 4: Legendre recovery probability as a function of sample size N : Top graphs are for noisy recovery, while bottom are for noiseless recovery. Horizontal Position corresponds to differing (d, p) .

uniform, and coherence-optimal measures. The colors represent the values of ρ_c , from which we observe the large gradient of ρ_c in terms of the random variables Ξ_1 along Ξ_2 . We consider a Hermite PC expansion of total order $p = 32$, giving $P = 561$ basis functions to approximate ρ_c . This high-order approximation is necessary due to the large gradient of ρ_c as evidenced by Figure 5 and the relatively slow decay of coefficients in the reference solution presented in Figure 6. This reference solution is computed using Gauss-Hermite quadrature approximation of the PC coefficients.

In Figure 7, we see plots of moments for the relative root-mean-squared

Figure 5: Realizations of the QoI $\rho_c = \rho(t = 4, \Xi_1, \Xi_2)$ corresponding to standard, uniform, and coherence-optimal samples of Ξ_1 and Ξ_2 .

Figure 6: PC coefficients of reference solution for the QoI $\rho_c = \rho(t = 4, \Xi_1, \Xi_2)$ in the surface reaction model. From left to right, coefficients are sorted by total-order of the polynomial.

error – between the reference and least squares solutions – as a function of the number of samples, N , obtained from 200 independent replications for each N . We find that the standard sampling fails to converge, while the asymptotic and coherence-optimal samplings lead to converged solution as N is increased. We note that, in all cases, the standard sampling resulted in rank-deficient matrices Ψ , and the reported results were generated using MATLAB's `lsqr` function with a large number of iterations.

Figure 7: Moments of root-mean-squared error between the reference and ℓ_2 -minimization solutions for the various sampling methods as a function of the number of samples.

4.4. Thermally Driven Cavity Flow

Following [2, 38, 39], we next consider a 2-D heat driven square cavity flow problem with uncertain wall temperature, as shown in Fig. 8. The left vertical wall has a deterministic, constant temperature $\tilde{T}_h$, referred to as the hot wall, while the right vertical wall has a stochastic temperature $\tilde{T}_c < \tilde{T}_h$ with constant mean $\tilde{\bar{T}}_c$, referred to as the cold wall. Both the top and bottom walls are assumed to be adiabatic. The reference temperature and the reference temperature difference are defined as $\tilde{T}_{ref} = (\tilde{T}_h + \tilde{\bar{T}}_c)/2$ and $\Delta\tilde{T}_{ref} = \tilde{T}_h - \tilde{\bar{T}}_c$, respectively. Under the assumption of small temperature differences, i.e., Boussinesq approximation, the governing equations in dimensionless variables are given by

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla p + \frac{\text{Pr}}{\sqrt{\text{Ra}}} \nabla^2 \mathbf{u} + \text{Pr} T \hat{\mathbf{y}}, \\ \nabla \cdot \mathbf{u} &= 0, \\ \frac{\partial T}{\partial t} + \nabla \cdot (\mathbf{u} T) &= \frac{1}{\sqrt{\text{Ra}}} \nabla^2 T, \end{aligned} \tag{39}$$

where $\hat{\mathbf{y}}$ is the unit vector $(0, 1)$, $\mathbf{u} = (u, v)$ is velocity vector field, $T = (\tilde{T} - \tilde{T}_{ref})/\Delta\tilde{T}_{ref}$ is normalized temperature ($\tilde{T}$ denotes non-dimensional temperature), p is pressure, and t is time. Non-dimensional Prandtl and Rayleigh

numbers are defined, respectively, as $\text{Pr} = \tilde{\mu}\tilde{c}_p/\tilde{\kappa}$ and $\text{Ra} = \tilde{\rho}g\beta\Delta\tilde{T}_{ref}\tilde{L}^3/(\tilde{\mu}\tilde{\kappa})$, where the superscript tilde ($\tilde{\cdot}$) denotes the non-dimensional quantities. Specifically, $\tilde{\mu}$ is molecular viscosity, $\tilde{\kappa}$ is thermal diffusivity, $\tilde{\rho}$ is density, g is gravitational acceleration, the coefficient of thermal expansion is given by β , and $\tilde{L}$ is the reference length. In this example, the Prandtl and Rayleigh numbers are set to $\text{Pr} = 0.71$ and $\text{Ra} = 10^6$, respectively. For more details on the non-dimensional variables in (39), we refer the reader to [2, Chapter 6.2]. On the cold wall, we use a (normalized) temperature distribution with stochastic fluctuations of the form

$$\begin{aligned} T_c(x=1, y, \Xi) &= \bar{T}_c + T'_c, \\ T'_c &= \sigma_T \sum_{i=1}^d \sqrt{\lambda_i} \varphi_i(y) \Xi_i, \end{aligned} \tag{40}$$

where $\bar{T}_c$ is a constant mean temperature. In (40), Ξ_i , $i = 1, \dots, d$, are independent random variables uniformly distributed on $[-1, 1]$. $\{\lambda_i\}_{i=1}^d$ and $\{\varphi_i(y)\}_{i=1}^d$ are the d largest eigenvalues and the corresponding eigenfunctions of the exponential covariance kernel $C_{T_c T_c}(y_1, y_2) = \exp\left(-\frac{|y_1 - y_2|}{l_c}\right)$, where l_c is the correlation length. In this setting, a (semi-)analytic representation of the eigenpairs $(\lambda_i, \varphi_i(y))$ in (40) is available, see, e.g., [1].

In our numerical tests, we let $(T_h, \bar{T}_c) = (0.5, -0.5)$, $d = 20$, $l_c = 1/21$, and $\sigma_T = 11/100$. Our QoI, the vertical velocity component at $(x, y) = (0.25, 0.25)$ denoted by $v(0.25, 0.25)$, is expanded in the Legendre PC basis of total degree $p = 4$ with only the first $P = 2500$ basis functions retained. We seek to accurately reconstruct $v(0.25, 0.25)$ with $N < P$ random samples of Ξ and the corresponding realizations of $v(0.25, 0.25)$. These samples are constructed from a least squares reference solution constructed from these 2500 basis polynomials as computed from a large number of Monte Carlo samples.

To identify $u(\Xi)$ as a function of the random inputs Ξ , we use a Legendre PC expansion of total order $p = 3$, which for this $d = 20$ stochastic dimensional problem yields $P = 1771$ basis functions. We note that the root-mean-squared error is considered here as the primary measure of recovery, and that this leaves a truncation error from the 729 basis functions not used in the recovery.

We investigate the ability to recover $u(\Xi)$ via (6), using each of the three sampling schemes considered for Legendre polynomials. For this problem

Figure 8: Illustration of the cavity flow problem.

we further improve the quality of the MCMC sampling through an initial burn-in of 1,000 discarded samples [35]. We provide bootstrapped estimates of the various moment based measures from a pool of samples generated beforehand. Specifically, samples for each realization are drawn from a pool of 50,000 previously generated samples, which are used to calculate bootstrap estimates of averages and standard deviations.

In Figure 9, we see plots of computed moments for the distribution of the relative root-mean-squared error between the computed and reference solutions obtained from 100 independent replications for each sample size, N . We note the standard and coherence-optimal sampling offer significant improvements over asymptotic, i.e., Chebyshev, sampling using similar sample sizes, N , both in terms of accuracy and robustness to differing realized samples. These observations are compatible with the theoretical results of Section 3 demonstrating a smaller coherence for the uniform sampling – as compared to Chebyshev sampling – for the case of $d > p$. The coherence-optimal sampling by construction leads to the theoretically smallest coherence, and provides a useful comparison for the other sampling schemes.

5. Proofs of Theorems in Section 2

We now present proofs for the theorems in Section 2, where we use the truncation of Ω to $\mathcal{S}$, and believe it is clear how to alter proofs to analogous

Figure 9: Plots for the moments of root-mean-squared error for independent residuals for the various sampling methods as a function of the number of samples.

results when this truncation is not made. We note that the recovery from the least squares problem (6) is connected with the spectral convergence of

$$\mathbf{M} := \frac{1}{N} (\mathbf{W}\Psi)^T \mathbf{W}\Psi \quad (41)$$

to the identity matrix, see, e.g., [18] for $\mathbf{W} = \mathbf{I}$. To insure stability, we wish to guarantee that the largest singular value of $\mathbf{M} - \mathbf{I}$, denoted by σ , is in the interval $[0, 1/2]$ with high probability. Indeed, we will consider any $\mathbf{M}$ that leads to $\sigma > 1/2$ as unstable, and consider the entire sampling a failure. As a result, we first wish to demonstrate that this failure probability is low, noting the immediate applicability from Theorem 1 of [18] when (16) is used. Throughout we use the restricted expectation notation as defined in (20), which is more convenient than the closely related conditional expectation. The following theorems also depend on a sampling event which we denote $\mathcal{E}$ that is the intersection of two events.

The first is that all N samples of $\mathbf{Y}$ would fall within $\mathcal{S}^c$, and consider any such realizations to lead to a failure for the entire sample. Using the definition of $\mu_{\mathcal{S}}$ in (17), we can guarantee that each realization only fails with probability at most $1/(NP)$, and so via the union bound, this failure probability is no larger than $1/P$, which is the value that is seen in our proofs. While this term is important when considering e.g. standard sampling of

Hermite polynomials where there is a probability that $\Xi \in \mathcal{S}^c$, practically, the sampling distribution for $\mathbf{Y}$ can be supported only $\mathcal{S}$, such as via a truncated normal distribution, and this term may be dropped. We note that while we use $1/P$, tight bounds or exact computations of this failure probability are possible, which will typically make this term far smaller.

The second event which we will require for $\mathcal{E}$ is that the matrix $\mathbf{M}$ is close enough to identity so that the spectral norm $\|\mathbf{M} - \mathbf{I}\| < 1/2$, an event which we will investigate more thoroughly.

We will use the condition,

$$\sum_{k=1}^P \mathbb{E} [|w(\mathbf{Y})\psi_k(\mathbf{Y})|^2 \mathbf{1}_{\mathcal{S}^c}] \leq \frac{1}{20\sqrt{P}},$$

to insure that $\mathcal{S}^c$ is not important in the ℓ_2 and, as we will now investigate, a spectral sense. The following lemma relates the convergence of $\mathbf{M}$ to $\mathbf{I}$ to a bound on the ℓ_2 -norm of the random rows. It is related to Lemma 1 of [18], and Section B of [22].

Lemma 5.1. *Let*

$$b := \sum_{k=1}^P \mathbb{E} [|w(\mathbf{Y})\psi_k(\mathbf{Y})|^2 \mathbf{1}_{\mathcal{S}^c}]$$

denote the bias in ℓ_2 -variation that is lost by restricting our sampled set to $\mathcal{S}$. That is, for

$$(\mathbf{w}\psi)_i := (w(\mathbf{Y}_i)\psi_1(\mathbf{Y}_i), \dots, w(\mathbf{Y}_i)\psi_P(\mathbf{Y}_i))^T,$$

it follows that

$$\mathbb{E} ((\mathbf{w}\psi)_i^T (\mathbf{w}\psi)_i; \mathcal{S}) = P - b.$$

If $\|\cdot\|$ denotes the spectral norm of the matrix, and $N\sqrt{P} \geq 20$, then

$$\mathbb{P} \left(\|\mathbf{M} - \mathbf{I}\| > 1/2 \mid \mathbf{Y}_i \in \mathcal{S} \ \forall i \right) \leq 2P \exp(-0.1N\mu_{\mathcal{S}}^{-1}).$$

Proof. We have that

$$\mathbf{I} = \mathbb{E} ((\mathbf{w}\psi)(\mathbf{w}\psi)^T | \mathcal{S}) \mathbb{P}(\mathcal{S}) + \mathbb{E} ((\mathbf{w}\psi)(\mathbf{w}\psi)^T | \mathcal{S}^c) \mathbb{P}(\mathcal{S}^c).$$

Letting $\|\cdot\|$ denote the spectral norm of the matrix. A brief calculation gives that

$$\|\mathbb{E}((\mathbf{w}\psi)(\mathbf{w}\psi)^T|\mathcal{S}) - \mathbf{I}\| \leq \frac{\mathbb{P}(\mathcal{S}^c)}{\mathbb{P}(\mathcal{S})} \left(\left\| \mathbb{E} \left((\mathbf{w}\psi)(\mathbf{w}\psi)^T \middle| \mathcal{S}^c \right) \right\| + 1 \right).$$

Using Jensen's inequality we have that,

$$\|\mathbb{E}((\mathbf{w}\psi)(\mathbf{w}\psi)^T|\mathcal{S}) - \mathbf{I}\| \leq \frac{\mathbb{P}(\mathcal{S}^c)}{\mathbb{P}(\mathcal{S})} \left[\mathbb{E} \left(\|\mathbf{w}\psi\|^2 \middle| \mathcal{S}^c \right) + 1 \right].$$

From the conditions in (17), this can be bounded by

$$\|\mathbb{E}((\mathbf{w}\psi)(\mathbf{w}\psi)^T|\mathcal{S}) - \mathbf{I}\| \leq \frac{1}{1 - (NP)^{-1}} \left[\frac{1}{20\sqrt{P}} + \frac{1}{NP} \right].$$

If $N\sqrt{P} \geq 20$, then this is bounded by

$$\|\mathbb{E}((\mathbf{w}\psi)(\mathbf{w}\psi)^T|\mathcal{S}) - \mathbf{I}\| \leq \frac{1}{10\sqrt{P}} =: \epsilon,$$

which should not be confused with a truncation error that also utilizes ϵ . We note that for this proof ϵ is not related to truncation error, but rather the bias in spectral convergence as a result of truncation of Ω to $\mathcal{S}$ in (17). This insures that if our samples come from $\mathcal{S}$, then the mean sample is nearly orthonormal, that is there is not much bias from this truncation. We utilize a Chernoff bound to show that there is a fast convergence to this mean. We have that

$$1 - \epsilon \leq \lambda_{\min}(\mathbb{E}(\mathbf{M})) \leq \lambda_{\max}(\mathbb{E}(\mathbf{M})) \leq 1 + \epsilon.$$

An application of the Chernoff bound as in Theorem 1.1 of [30] and Theorem 1 of [18] gives that for $\delta \in [0, 1]$

$$\begin{aligned} \mathbb{P} \left(\lambda_{\min}(\mathbf{M}) \leq (1 - \delta)(1 - \epsilon) \middle| \mathbf{Y}_i \in \mathcal{S} \ \forall i \right) &\leq P \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^{N(1-\epsilon)\mu_{\mathcal{S}}^{-1}}; \\ \mathbb{P} \left(\lambda_{\max}(\mathbf{M}) \geq (1 + \delta)(1 + \epsilon) \middle| \mathbf{Y}_i \in \mathcal{S} \ \forall i \right) &\leq P \left(\frac{e^{\delta}}{(1 + \delta)^{1+\delta}} \right)^{N(1+\epsilon)\mu_{\mathcal{S}}^{-1}}. \end{aligned}$$

Note that

$$\begin{aligned} (1 - \delta)(1 - \epsilon) \leq 1/2 &\implies \delta \geq \frac{1 - 2\epsilon}{2 - 2\epsilon}; \\ (1 + \delta)(1 + \epsilon) \geq 3/2 &\implies \delta \geq \frac{1 - 2\epsilon}{2 + 2\epsilon}. \end{aligned}$$

and so we have that a critical δ , denoted δ_* , such that for all $\delta < \delta_*$ the matrix $\mathbf{M}$ is guaranteed to satisfy $\|\mathbf{M} - \mathbf{I}\|_2 \leq 1/2$ is $\delta_* := (1 - 2\epsilon)/(2 + 2\epsilon)$ as it is the smaller tolerance. Note that for $0 \leq \delta < 1$, [18],

$$\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \geq \frac{e^{\delta}}{(1 + \delta)^{1+\delta}},$$

and so we may bound the sum of the probabilities by

$$\mathbb{P}\left(\|\mathbf{M} - \mathbf{I}\| > \frac{1}{2} \middle| \mathbf{Y}_i \in \mathcal{S} \forall i\right) \leq 2P \left(\frac{e^{-\delta_*}}{(1 - \delta_*)^{1-\delta_*}}\right)^{N(1-\epsilon)\mu_S^{-1}}.$$

We now bound this probability. We note that if $\epsilon \leq 1/2$ then

$$\delta_* \leq \frac{1 - 2\epsilon}{2 + 2\epsilon} \leq \frac{1}{2} - \epsilon.$$

To create a useful bound in terms of ϵ , we note that

$$\begin{aligned} c_\epsilon &:= 1/2 - \epsilon + (1/2 + \epsilon) \log(1/2 + \epsilon); \\ \frac{e^{-\delta_*}}{(1 - \delta_*)^{1-\delta_*}} &\leq \exp(-c_\epsilon). \end{aligned}$$

Thus, noting that if $P = 1$, then $\epsilon = 1/20$ and $c_\epsilon \approx 0.1212$,

$$\begin{aligned} \mathbb{P}\left(\|\mathbf{M} - \mathbf{I}\| > \frac{1}{2} \middle| \mathbf{Y}_i \in \mathcal{S} \forall i\right) &\leq 2P \exp(-c_\epsilon(1 - \epsilon)N\mu_S^{-1}); \\ &\leq 2P \exp(-0.95c_\epsilon N\mu_S^{-1}); \\ &\leq 2P \exp(-0.114N\mu_S^{-1}). \end{aligned}$$

As $P \rightarrow \infty$, $c_\epsilon \rightarrow 0.1534$, and so to does the multiplicative constant which is seen to not heavily depend on P . $\square$

This Lemma shows that with a probability that converges to 1 exponentially in $N/\mu_{\mathcal{S}}$, $\mathbf{M}$ is well-conditioned. Changes could allow us to lower the bound on $\|\mathbf{M} - \mathbf{I}\|$ from 1/2 at the cost of a higher failure probability. This reduction on $\|\mathbf{M} - \mathbf{I}\|$ will be useful in a way that we see in the next proof. Also we could alter the conditions of (17) for $\mathcal{S}$ to affect this probability. We consider neither point here, but we refer the interested reader to Theorem 1 of [18] and Section B of [22].

We are now equipped to prove Theorem 2.1 and 2.3, but before doing so we first briefly show that the truncation error that is orthogonal to the basis functions restricted to the smaller set $\mathcal{S}$ is less than that over the full set Ω .

Lemma 5.2. *Let $\epsilon_{\mathcal{S}}(\boldsymbol{\xi})$ be the truncated error that is orthogonal to the basis functions restricted to $\mathcal{S}$, and let $\epsilon_{\Omega}(\boldsymbol{\xi})$ be the error that is orthogonal to the basis functions over the full Ω . Then*

$$\mathbb{E}(\epsilon_{\mathcal{S}}^2(\boldsymbol{\Xi}); \mathcal{S}) \leq \mathbb{E}(\epsilon_{\Omega}^2(\boldsymbol{\Xi})).$$

We use this inequality in the proofs of Theorem 2.1 and 2.3, as our errors are based on the truncated orthogonal version, and we would like to substitute for the larger, more global error. We note that this theorem applies only to the theoretically optimal solutions, and not any computed approximation.

Proof. Let $\mathcal{V}$ be the span of $\{\psi_1(\boldsymbol{\xi}), \dots, \psi_P(\boldsymbol{\xi})\}$. We consider the least squares solution over $\mathcal{S}$, and denote the $L_2(\mathcal{S}, f)$ -optimal approximation to u as

$$\hat{u}_{\mathcal{S}} := \arg \min_{u_{\star} \in \mathcal{V}} \int_{\mathcal{S}} |u(\boldsymbol{\xi}) - u_{\star}(\boldsymbol{\xi})|^2 f(\boldsymbol{\xi}) d\boldsymbol{\xi}.$$

Similarly, we let $\hat{u}_{\Omega}$ denote the least squares solution that minimizes the above over Ω . It follows that

$$\begin{aligned} \int_{\mathcal{S}} |u(\boldsymbol{\xi}) - \hat{u}_{\mathcal{S}}(\boldsymbol{\xi})|^2 f(\boldsymbol{\xi}) d\boldsymbol{\xi} &\leq \int_{\mathcal{S}} |u(\boldsymbol{\xi}) - \hat{u}_{\Omega}(\boldsymbol{\xi})|^2 f(\boldsymbol{\xi}) d\boldsymbol{\xi}, \\ &\leq \int_{\Omega} |u(\boldsymbol{\xi}) - \hat{u}_{\Omega}(\boldsymbol{\xi})|^2 f(\boldsymbol{\xi}) d\boldsymbol{\xi}. \end{aligned}$$

□

5.1. Proof of Theorem 2.1

We recall that our error is conditioned on a set $\mathcal{E}$ which is the intersection of two events. The first event, $\mathcal{E}_1$, is that all sampled $\mathbf{Y}$ are in the set $\mathcal{S}$, which may be a proper subset of Ω . The second event, $\mathcal{E}_2$, is that $\|\mathbf{M} - \mathbf{I}\| \leq 1/2$. We note that the probability density function $f(\boldsymbol{\xi})$ for $\boldsymbol{\Xi}$ is related to that for $\mathbf{Y}$ multiplied by $w^2(\boldsymbol{\xi})$. We are premultiplying by $w(\mathbf{Y})$ because of sampling $\mathbf{Y}$ instead of $\boldsymbol{\Xi}$, and so we are able to write expectations of $w^2(\mathbf{Y})g(\mathbf{Y})$ in terms of $\boldsymbol{\Xi}$. Further, for any appropriately measurable function g and appropriately measurable set $\mathcal{T}$,

$$\mathbb{E}(w^2(\mathbf{Y})g(\mathbf{Y}); \mathcal{T}) = \mathbb{E}(g(\boldsymbol{\Xi}); \boldsymbol{\Xi} \in \mathcal{T}).$$

As the event $\mathcal{E}$ includes that all samples are in $\mathcal{S}$, we thus have that

$$\mathbb{E}(w^2(\mathbf{Y})g(\mathbf{Y}); \mathcal{E}) = \mathbb{E}(g(\boldsymbol{\Xi}); \mathcal{E}).$$

The remainder of the proof borrows methodology from that of Theorem 2 of [18], while having several differences regarding the truncation of Ω to $\mathcal{S}$. Recalling that $\epsilon(\boldsymbol{\Xi})$ is the truncation error, we have that $w^2(\mathbf{Y})\epsilon(\mathbf{Y})$ is orthogonal to the space $\mathcal{V} = \text{span}\{\psi_1(\boldsymbol{\Xi}), \dots, \psi_P(\boldsymbol{\Xi})\}$. We may define $\hat{\mathbf{a}}$ to be the difference between the true theoretical solution and our computed solution by

$$\hat{\mathbf{a}} := \mathbf{c} - \hat{\mathbf{c}}$$

and note that this implies that

$$\mathbf{M}\hat{\mathbf{a}} = \mathbf{b},$$

where, as the rows of $\mathbf{u}$ and $\boldsymbol{\Psi}$ are each multiplied by $w(\mathbf{Y}_i)$ and $\hat{\mathbf{c}}$ is a least squares solution,

$$b_k = \frac{1}{N} \sum_{i=1}^N w^2(\mathbf{Y}_i) \epsilon(\mathbf{Y}_i) \psi_k(\mathbf{Y}_i).$$

As we assume the hypotheses of Lemma 5.1, we have that $\|\mathbf{M}^{-1}\| \leq 2$, and thus

$$\|\hat{\mathbf{a}}\|_2^2 \leq 4\|\mathbf{b}\|_2^2.$$

Applying the expectation, integral, and a triangular inequality gives that

$$\mathbb{E} \left(\int_{\Omega} (u(\boldsymbol{\xi}) - \hat{u}(\boldsymbol{\xi}))^2 f(\boldsymbol{\xi}) d\boldsymbol{\xi}; \mathcal{E} \right) \leq \mathbb{E}(\epsilon(\boldsymbol{\Xi})^2) + 4\mathbb{E}(\|\mathbf{b}\|_2^2; \mathcal{E}).$$

We now work to bound $\mathbb{E}(\|\mathbf{b}\|_2^2; \mathcal{E})$, noting that this includes that all draws are restricted to come from $\mathcal{S}$. We have that

$$\begin{aligned} \mathbb{E}(b_k^2; \mathcal{E}) &\leq \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \mathbb{E}(w^2(\mathbf{Y}_i) \epsilon(\mathbf{Y}_i) \psi_k(\mathbf{Y}_i) w^2(\mathbf{Y}_j) \epsilon(\mathbf{Y}_j) \psi_k(\mathbf{Y}_j); \mathcal{S}), \\ &\leq \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}(w^4(\mathbf{Y}_i) \epsilon^2(\mathbf{Y}_i) \psi_k^2(\mathbf{Y}_i); \mathcal{S}), \\ &\leq \frac{1}{N} \mathbb{E}(w^2(\mathbf{Y}) \epsilon^2(\mathbf{Y}); \mathcal{S}) \sup_{\boldsymbol{\xi} \in \mathcal{S}} |w^2(\boldsymbol{\xi}) \psi_k^2(\boldsymbol{\xi})|, \end{aligned}$$

where the first inequality comes from the definition of b_k ; the second from the independence of the $\mathbf{Y}_k$ and orthogonality of $\{\psi_k\}_{k=1}^P \cup \epsilon$; and the third is a direct point-wise bound. Hence, summing for $k = 1 : P$ and recalling the definition in (17), we arrive at

$$\begin{aligned} \mathbb{E}(\|\mathbf{b}\|_2^2; \mathcal{E}) &\leq \frac{\mathbb{E}(\epsilon^2(\boldsymbol{\Xi}))}{N} \cdot \sup_{\boldsymbol{\xi} \in \mathcal{S}} \sum_{k=1}^P |w(\boldsymbol{\xi}) \psi_k(\boldsymbol{\xi})|^2, \\ &\leq \frac{\mathbb{E}(\epsilon^2(\boldsymbol{\Xi}))}{N} \cdot \mu_{\mathcal{S}}. \end{aligned}$$

Putting the above pieces together, we have that

$$\mathbb{E} \left(\int_{\Omega} (u(\boldsymbol{\xi}) - \hat{u}(\boldsymbol{\xi}))^2 f(\boldsymbol{\xi}) d\boldsymbol{\xi}; \mathcal{E} \right) \leq \mathbb{E}(\epsilon^2(\boldsymbol{\Xi})) \left(1 + \frac{4\mu_{\mathcal{S}}}{N} \right).$$

This completes the proof.

5.2. Proof of Theorem 2.3

This proof is related to Theorem 3 of [18], although we bound the measurement error similarly to the bound on truncation error. We let ϵ_T denote the truncation error (denoted by ϵ in the previous theorem), with corresponding hat notation. Additionally we let $\epsilon_M(\mathbf{Y}) := \mathbb{E}(\epsilon_M(\boldsymbol{\chi}; \mathbf{Y}) | \mathbf{Y})$, be itself a random quantity that represents the contribution arising from the

unobserved random vector, $\boldsymbol{\chi}$. Similarly to before, we define $\hat{\mathbf{a}}_m$ to be the true solution, minus the computed solution as well as the contribution defined above that arises from error due to the truncation of the basis, leaving only a solution computation associated with measurement error. That is,

$$\hat{\mathbf{a}}_m := \mathbf{c} - \hat{\mathbf{c}} - \hat{\mathbf{a}}$$

is the portion of the solution which fits the measurement noise and satisfies

$$\mathbf{M}\hat{\mathbf{a}}_m = \mathbf{d}.$$

Here, $\mathbf{d}$ is given by

$$d_k = \frac{1}{N} \sum_{i=1}^N w^2(\mathbf{Y}_i) \epsilon_M(\mathbf{Y}_i) \psi_k(\mathbf{Y}_i).$$

Applying the integral and expectation similarly to before we have that,

$$\mathbb{E} \left(\int_{\Omega} (u(\boldsymbol{\xi}) - \hat{u}(\boldsymbol{\xi}))^2 f(\boldsymbol{\xi}) d\boldsymbol{\xi}; \mathcal{E} \right) \leq \mathbb{E}(\epsilon^2(\boldsymbol{\Xi})) + 4 \cdot [2\mathbb{E}(\|\mathbf{b}\|_2^2; \mathcal{E}) + 2\mathbb{E}(\|\mathbf{d}\|_2^2; \mathcal{E})]. \quad (42)$$

We bound $\mathbb{E}(d_k^2; \mathcal{E})$ similarly to how we bounded $\mathbb{E}(b_k^2; \mathcal{E})$. Particularly,

$$\begin{aligned} \mathbb{E}(d_k^2; \mathcal{E}) &\leq \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \mathbb{E}(w^2(\mathbf{Y}_i) \epsilon_M(\mathbf{Y}_i) \psi_k(\mathbf{Y}_i) w^2(\mathbf{Y}_j) \epsilon_M(\mathbf{Y}_j); \mathcal{S}), \\ &\leq \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}(w^4(\mathbf{Y}_i) \epsilon_M^2(\mathbf{Y}_i) \psi_k^2(\mathbf{Y}_i); \mathcal{S}), \\ &\leq \frac{1}{N} \mathbb{E}(w^2(\mathbf{Y}) \epsilon_M^2(\mathbf{Y}); \mathcal{S}) \sup_{\boldsymbol{\xi} \in \mathcal{S}} |w^2(\boldsymbol{\xi}) \psi_k^2(\boldsymbol{\xi})|. \end{aligned}$$

Which gives that

$$\mathbb{E}(\|\mathbf{d}\|_2^2; \mathcal{E}) \leq \frac{\mathbb{E}(\epsilon_M^2(\boldsymbol{\Xi})) \mu_{\mathcal{S}}}{N}.$$

Using this and the results from Theorem 2.1 plugged into (42), completes the proof.

5.3. Proof of Theorem 3.3

While the first part of this follows from the proof of Theorem 4.5 of [23], we briefly prove the second statement, which arises from $B^2(\boldsymbol{\xi})$ being a linear combination of squared basis functions. We note that if we take $w(\boldsymbol{\xi})$ to be $1/B(\boldsymbol{\xi})$ then

$$\begin{aligned}\mu_{\Omega}(\mathbf{Y}) &:= \max_{\boldsymbol{\xi} \in \Omega} c^{-2} w^2(\boldsymbol{\xi}) \sum_{k=1}^P \psi_k^2(\boldsymbol{\xi}), \\ &= \max_{\boldsymbol{\xi} \in \Omega} c^{-2} = c^{-2},\end{aligned}$$

where it remains to identify the normalization constant c^{-2} so that $f_{\mathbf{Y}}(\boldsymbol{\xi}) = c^2 f(\boldsymbol{\xi})/w^2(\boldsymbol{\xi})$ is a probability distribution. Recalling the definition of w in this case we have that

$$\begin{aligned}1 &= \int_{\boldsymbol{\xi} \in \Omega} f_{\mathbf{Y}}(\boldsymbol{\xi}) d\boldsymbol{\xi}, \\ &= c^2 \sum_{k=1}^P \int_{\boldsymbol{\xi} \in \Omega} f(\boldsymbol{\xi}) \psi_k^2(\boldsymbol{\xi}) d\boldsymbol{\xi}.\end{aligned}$$

By the orthonormality of the $\{\psi_k(\boldsymbol{\xi})\}_{k=1}^P$ it follows that $c^2 = 1/P$, implying that $\mu_{\Omega} = c^{-2} = P$.

6. Conclusions

We provided an analysis of Hermite and Legendre polynomials which allowed us to bound a coherence parameter and generate recovery guarantees for polynomial chaos expansions obtained via least squares regression. We also identified alternative random sampling schemes which provide sharper guarantees and demonstrate improved polynomial chaos reconstructions relative to the random sampling from the orthogonality measure of these bases. These sampling methods were derived based on the properties of Hermite and Legendre polynomials. Furthermore, we showed a Markov Chain Monte Carlo method for generating samples that minimizes the coherence parameter, thereby achieving an optimality for the number of random solution realizations. Such a sampling was referred to as coherence-optimal sampling, and guarantees recovery with a number of samples that scales linearly in the number of basis functions, up to a **logarithmic factor**. Positive results

were observed when computing the solution to a non-linear ordinary differential equation, where a high order Hermite polynomial chaos expansion was needed for an accurate solution approximation. Similarly, sampling methods were compared on arbitrary manufactured stochastic functions, and the different sampling strategies were tested for identifying the solution of a 20-dimensional heat driven fluid flow. In all examples, positive results were obtained for the coherence-optimal sampling method.

Acknowledgements

We thank Ji Peng from CU Boulder for providing the deterministic codes for the problem in Section 4.4.

This material is based upon work supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, under Award Number DE-SC0006402, as well as National Science Foundation under grant DMS-1228359.

References

- [1] R. Ghanem, P. Spanos, Stochastic Finite Elements: A Spectral Approach, Springer Verlag, 1991.
- [2] O. L. Maitre, O. Knio, Spectral Methods for Uncertainty Quantification with Applications to Computational Fluid Dynamics, Springer, 2010.
- [3] D. Xiu, Numerical Methods for Stochastic Computations: A Spectral Method Approach, Princeton University Press, 2010.
- [4] R. Smith, Uncertainty Quantification: Theory, Implementation, and Applications, Vol. 12, SIAM, 2013.
- [5] D. Xiu, G. Karniadakis, The Wiener-Askey polynomial chaos for stochastic differential equations, SIAM Journal on Scientific Computing 24 (2) (2002) 619–644.
- [6] C. Soize, R. Ghanem, Physical systems with random uncertainties: Chaos representations with arbitrary probability measure, SIAM Journal of Scientific Computing 26 (2) (2005) 395–410.

- [7] I. Babuška, F. Nobile, R. Tempone, A stochastic collocation method for elliptic partial differential equations with random input data, *SIAM J. Numer. Anal.* 45 (3) (2007) 1005–1034.
- [8] L. Mathelin, M. Hussaini, A stochastic collocation algorithm for uncertainty analysis, Tech. Rep. NAS 1.26:212153; NASA/CR-2003-212153, NASA Langley Research Center (2003).
- [9] H. Li, D. Zhang, Probabilistic collocation method for flow in porous media: Comparisons with other stochastic methods, *Water Resources Research* 43 (9).
- [10] B. Ganis, H. Klie, M. F. Wheeler, T. Wildey, I. Yotov, D. Zhang, Stochastic collocation and mixed finite elements for flow in porous media, *Computer Methods in Applied Mechanics and Engineering* 197 (4344) (2008) 3547 – 3559, stochastic Modeling of Multiscale and Multiphysics Problems.
- [11] S. Hosder, R. Walters, 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.
- [12] M. Berveiller, B. Sudret, M. Lemaire, Stochastic finite element: a non intrusive approach by regression, *European Journal of Computational Mechanics/Revue Européenne de Mécanique Numérique* 15 (1-3) (2006) 81–92.
- [13] B. Jones, A. Doostan, G. Born, Nonlinear propagation of orbit uncertainty using non-intrusive polynomial chaos, *Journal of Guidance, Control, and Dynamics* 36 (2) (2013) 430–444.
- [14] B. Jones, A. Doostan, Satellite collision probability estimation using polynomial chaos expansions, *Advances in Space Research* 52 (11) (2013) 1860–1875.
- [15] Z. Gao, T. Zhou, Choice of nodal sets for least square polynomial chaos method with application to uncertainty quantification, *Commun. Comput. Phys* (2014) 365–381.

- [16] G. Migliorati, F. Nobile, E. Schwerin, R. Tempone, Analysis of discrete L^2 projection on polynomial spaces with random evaluations, *Foundations of Computational Mathematics* 14 (3) (2014) 419–456.
- [17] G. Migliorati, F. Nobile, E. von Schwerin, R. Tempone, Approximation of quantities of interest in stochastic pdes by the random discrete L^2 projection on polynomial spaces, *SIAM Journal on Scientific Computing* 35 (3) (2013) A1440–A1460.
- [18] A. Cohen, M. A. Davenport, D. Leviatan, On the stability and accuracy of least squares approximations., *Foundations of Computational Mathematics* 13 (5) (2013) 819–834.
- [19] A. Chkifa, A. Cohen, G. Migliorati, F. Nobile, R. Tempone, Discrete least squares polynomial approximation with random evaluations; application to parametric and stochastic elliptic pdes, Also available as EPFL-MATHICSE report (2013) 35–2013.
- [20] T. Zhou, A. Narayan, Z. Xu, Multivariate discrete least-squares approximations with a new type of collocation grid, *arXiv preprint arXiv:1401.0894*.
- [21] T. Tang, T. Zhou, On discrete least square projection in unbounded domain with random evaluations and its application to parametric uncertainty quantification, *arXiv preprint arXiv:1403.6579*.
- [22] E. J. Candés, Y. Plan, A probabilistic and ripless theory of compressed sensing, *Information Theory, IEEE Transactions on* 57 (11) (2010) 7235–7254.
- [23] J. Hampton, A. Doostan, Compressive sampling of polynomial chaos expansions: Convergence analysis and sampling strategies, *Journal of Computational Physics* Under revision.
- [24] H. Rauhut, R. Ward, Sparse legendre expansions via l1-minimization, *J. Approx. Theory* 164 (5) (2012) 517–533.
- [25] G. Szegő, *Orthogonal Polynomials*, American Mathematical Society, American Mathematical Society, 1939.

- [26] D. Dominici, Asymptotic analysis of the hermite polynomials from their differential-difference equation, *Journal of Difference Equations and Applications* 13 (12) (2007) 1115–1128.
- [27] R. Askey, S. Wainger, Mean convergence of expansions in laguerre and hermite series, *American Journal of Mathematics* 87 (3) (1965) 695–708.
- [28] B. Muckenhoupt, Asymptotic forms for laguerre polynomials, *Proceedings of the American Mathematical Society* 24 (2) (1970) 288–292.
- [29] P. Nevai, T. Erdélyi, A. Magnus, Generalized jacobi weights, christoffel functions, and jacobi polynomials, *SIAM Journal on Mathematical Analysis* 25 (2) (1994) 602–614.
- [30] J. A. Tropp, User-friendly tail bounds for sums of random matrices, *Foundations of Computational Mathematics* 12 (4) (2012) 389–434.
- [31] A. Doostan, H. Owhadi, A non-adapted sparse approximation of PDEs with stochastic inputs, *Journal of Computational Physics* 230 (2011) 3015–3034.
- [32] L. Mathelin, K. Gallivan, A compressed sensing approach for partial differential equations with random input data, *Commun. Comput. Phys.* 12 (2012) 919–954.
- [33] L. Yan, L. Guo, D. Xiu, Stochastic collocation algorithms using ℓ_1 -minimization, *International Journal for Uncertainty Quantification* 2 (3).
- [34] W. K. HASTINGS, Monte carlo sampling methods using markov chains and their applications, *Biometrika* 57 (1) (1970) 97–109.
- [35] W. R. Gilks, S. Richardson, D. J. Spiegelhalter, Markov chain Monte Carlo in practice, Vol. 2, CRC press, 1996.
- [36] A. Makeev, D. Maroudas, I. Kevrekidis, Coarse stability and bifurcation analysis using stochastic simulators: Kinetic Monte Carlo examples, *The Journal of chemical physics* 116 (23) (2002) 10083–10091.
- [37] O. L. Maître, H. Najm, R. Ghanem, O. Knio, Multi-resolution analysis of wiener-type uncertainty propagation schemes, *J. Comput. Phys* 197 (2004) 502–531.

- 1
2
3
4
5
6
7
8
9 [38] P. L. Quéré, Accurate solutions to the square thermally driven cavity at
10 high rayleigh number, *Computers & Fluids* 20 (1) (1991) 29–41.
11
12 [39] J. Peng, J. Hampton, A. Doostan, A weighted ℓ_1 -minimization ap-
13 proach for sparse polynomial chaos expansions, *Journal of Computa-*
14 tional Physics 267 (2014) 92–111.
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65