On polynomial chaos expansion via gradient-enhanced ℓ_1-minimization

Ji Peng[a], Jerrad Hampton[b], Alireza Doostan[b],*

[a] *Mechanical Engineering Department, University of Colorado, Boulder, CO 80309, USA*
[b] *Aerospace Engineering Sciences Department, University of Colorado, Boulder, CO 80309, USA*

A B S T R A C T

Gradient-enhanced Uncertainty Quantification (UQ) has received recent attention, in which the derivatives of a Quantity of Interest (QoI) with respect to the uncertain parameters are utilized to improve the surrogate approximation. Polynomial chaos expansions (PCEs) are often employed in UQ, and when the QoI can be represented by a sparse PCE, ℓ_1-minimization can identify the PCE coefficients with a relatively small number of samples. In this work, we investigate a gradient-enhanced ℓ_1-minimization, where derivative information is computed to accelerate the identification of the PCE coefficients. For this approach, stability and convergence analysis are lacking, and thus we address these here with a probabilistic result. In particular, with an appropriate normalization, we show the inclusion of derivative information will almost-surely lead to improved conditions, e.g. related to the null-space and coherence of the measurement matrix, for a successful solution recovery. Further, we demonstrate our analysis empirically via three numerical examples: a manufactured PCE, an elliptic partial differential equation with random inputs, and a plane Poiseuille flow with random boundaries. These examples all suggest that including derivative information admits solution recovery at reduced computational cost.

© 2015 Elsevier Inc. All rights reserved.

1. Introduction

In complex engineering system analysis, inherent variability in inputs and imperfect knowledge of the physics can lead to unfounded confidence or unnecessary diffidence in understanding Quantities of Interest (QoI). Uncertainty Quantification (UQ) [1–3] as a study aims to develop numerical tools that can accurately predict QoI and facilitate the quantitative validation of the simulation model.

To characterize uncertainty, probability is a natural framework. We model the uncertain inputs as a d-dimensional vector of independent random variables $\Xi := (\Xi_1, \ldots, \Xi_d)$, with probability density function $\rho(\Xi)$. The QoI that we seek to approximate is denoted by a scalar $u(\Xi)$. Here we utilize polynomial chaos expansions (PCEs) [1,4] to approximate $u(\Xi)$, assumed to have finite variance. In this case, $u(\Xi)$ can be represented as an expansion in multivariate orthogonal polynomials $\psi_j(\Xi)$, i.e.,

---

* Corresponding author.
  E-mail address: alireza.doostan@colorado.edu (A. Doostan).

$$u(\mathbf{\Xi}) = \sum_{j=1}^{\infty} c_j \psi_j(\mathbf{\Xi}) \approx \sum_{j=1}^{P} c_j \psi_j(\mathbf{\Xi}) + \epsilon_t(\mathbf{\Xi}), \tag{1}$$

where c_j, $j = 1, 2, \ldots$, are the corresponding PCE coefficients, and ϵ_t is the truncation error associated with retaining P terms of a sorted basis. The PCE coefficients can be computed by the projection

$$c_j = \int u(\mathbf{\Xi}) \psi_j(\mathbf{\Xi}) \rho(\mathbf{\Xi}) d\mathbf{\Xi} = \mathbb{E}\left[u(\mathbf{\Xi}) \psi_j(\mathbf{\Xi})\right], \tag{2}$$

where the operator $\mathbb{E}$ denotes the mathematical expectation. Here we assume that the $\psi_j(\mathbf{\Xi})$ are normalized such that $\mathbb{E}\left[\psi_j^2(\mathbf{\Xi})\right] = 1$.

Typically, a full P term approximation is not necessary, and we can restrict to an unknown subset $\mathcal{C} \subset \{1, \ldots, P\}$ such that $|\mathcal{C}|$, the number of elements in $\mathcal{C}$, is significantly smaller than P. In addition, $|\mathcal{C}|$ is referred to as the sparsity of the approximation. We approximate $u(\mathbf{\Xi})$ in (1) then by

$$u(\mathbf{\Xi}) \approx \sum_{j \in \mathcal{C}} c_j \psi_j(\mathbf{\Xi}). \tag{3}$$

This concept of using one basis, indexed by $\{1, \ldots, P\}$, to compute an approximation in a small but unknown subset of those basis functions, indexed by $\mathcal{C}$, falls within the context of compressive sampling [5–9].

To identify PCE coefficients, we consider non-intrusive sampling methods, in which deterministic solvers for the QoI are not modified. Such methods include Monte Carlo simulation [10,2], pseudo-spectral stochastic collocation [11,12,2,13], least squares regression [14–16], and ℓ_1-minimization [17,7,18,19,8,9,20,21]. In this work, we adopt ℓ_1-minimization to estimate the coefficients, solving the problem

$$\arg\min_{\mathbf{c}} \|\mathbf{c}\|_1 \quad \text{subject to} \quad \|\mathbf{u} - \mathbf{\Psi}\mathbf{c}\|_2 \leq \delta, \tag{4}$$

where the vector $\mathbf{c} := (c_1, \ldots, c_P)^T$ contains PCE coefficients, while the vector $\mathbf{u}$ and the so-called *measurement matrix* $\mathbf{\Psi}$ contain function evaluations at realizations of the random input, $\mathbf{\Xi}$. Specifically, denoting the ith realization of $\mathbf{\Xi}$ as $\boldsymbol{\xi}^{(i)}$,

$$\mathbf{u} := \left(u(\boldsymbol{\xi}^{(1)}), \ldots, u(\boldsymbol{\xi}^{(N)})\right)^T; \tag{5}$$

$$\mathbf{\Psi}(i, j) := \psi_j\left(\boldsymbol{\xi}^{(i)}\right). \tag{6}$$

In (4), δ is a tolerance parameter necessitated by the truncation error, $\epsilon_t(\mathbf{\Xi})$, to reduce the effect of overfitting. For this work, δ is identified via cross-validation [22].

It has been shown that as the number of samples, N, increases, the probability of accurately recovering PCE coefficients experiences a corresponding increase [9]. Often the deterministic solver of u is computationally demanding, and it is expensive to compute realized $u(\boldsymbol{\xi}^{(i)})$. In order to deal with this situation, coefficient estimation based on gradient-enhanced PCE has received recent attention [23–27,21]. Specifically, the gradient information utilized is

$$\mathbf{u}_\partial := \left(\frac{\partial u}{\partial \Xi_1}\left(\boldsymbol{\xi}^{(1)}\right), \ldots, \frac{\partial u}{\partial \Xi_d}\left(\boldsymbol{\xi}^{(1)}\right), \ldots, \frac{\partial u}{\partial \Xi_1}\left(\boldsymbol{\xi}^{(N)}\right), \ldots, \frac{\partial u}{\partial \Xi_d}\left(\boldsymbol{\xi}^{(N)}\right)\right)^T \tag{7}$$

and

$$\mathbf{\Psi}_\partial((i-1) \cdot d + k, j) := \frac{\partial \psi_j}{\partial \Xi_k}(\boldsymbol{\xi}^{(i)}), \quad k = 1, \ldots, d. \tag{8}$$

With this gradient-enhancement, $\mathbf{c}$ is solved by minimizing $\|\mathbf{c}\|_2$, in least squares regression, or $\|\mathbf{c}\|_1$, in compressive sampling, subject to a constraint on

$$\left\|\begin{pmatrix} \mathbf{u} \\ \mathbf{u}_\partial \end{pmatrix} - \begin{pmatrix} \mathbf{\Psi} \\ \mathbf{\Psi}_\partial \end{pmatrix} \mathbf{W} \mathbf{c}\right\|_2, \tag{9}$$

where $\mathbf{W}$ is a positive-diagonal matrix depending on the basis functions and their derivatives, which we shall specify in Section 3.2. We note that these weights are chosen to maintain orthogonality (in expectation) in the matrix of basis evaluations, and that in practice it may be preferable to alter these weights to incorporate the normalizations given in [28] or [29]. In this way, N realizations of the random inputs $\mathbf{\Xi}$ provide an $N(d+1) \times P$ matrix, where the gradients $\partial u/\partial \Xi_k$, $k = 1, \ldots, d$, may be computed at each evaluation of $\mathbf{\Xi}$ from, e.g., direct or adjoint sensitivity equations [30], or automatic differentiation [31]. We assume that $u(\mathbf{\Xi})$ and its partial derivatives are square integrable with respect to the measure for $\mathbf{\Xi}$, and are represented in the appropriate basis such that the QoI and its derivatives correspond to identical coefficients.

In a related context, a gradient-enhanced sparse approximation based on ℓ_1-minimization was proposed in [32], in which derivatives $\partial u/\partial \Xi_k$, $k = 1, \ldots, d$, are projected onto Legendre PCE, with a corresponding coefficient vector $\mathbf{c}_\partial$.

1.1. Contribution of this work

In this work, we investigate a gradient-enhanced ℓ_1-minimization approach as seen in [21], in which the derivative information is empirically shown to improve the resolution of PCE coefficients.

In Section 3, a theoretical contribution concerning the Restricted Isometry Constant (RIC) is presented regarding recovery in ℓ_1-minimization with derivative information based on Hermite PCE. The RIC is a constant associated with the measurement matrix that provides fruitful (probabilistic) bounds for recovery of solutions via ℓ_1-minimization [33]. Additionally, under an appropriate normalization introduced by $\boldsymbol{W}$ in (9), it is guaranteed that null-space, measurement matrix column inner-products, and coherence related measures are almost-surely improved, implying that stability for solutions computed by (4) are not reduced by including derivative information. These analyses are original to the authors' best knowledge, and provide a framework for analysis of recovery in other PCE bases. Also, though not considered here, the approach may be extended to least squares polynomial chaos regression [16].

Three numerical experiments are used to demonstrate the gradient-enhanced ℓ_1-minimization in Section 4, among which a plane Poiseuille flow with random boundaries is simulated, and the derivative information is approximated via an adjoint sensitivity method. The empirical results agree with the theoretical analysis, and suggest that the inclusion of derivative information can improve solution recovery at a lower overall computational cost.

The structure of the manuscript is as follows. In Section 2, we state our problem and introduce the formulation of the gradient-enhanced ℓ_1-minimization approach. In Section 3, we present theoretical results concerning the stability and convergence of the gradient-enhanced ℓ_1-minimization approach, in the context of Hermite PCE. In Section 4, we demonstrate our analysis empirically via three numerical examples: a manufactured PCE, an elliptic partial differential equation with random inputs, and a plane Poiseuille flow with random boundaries. Section 5 presents the proofs to the results in Section 3.

2. Method synopsis

2.1. Problem statement

We use differential equations to model engineering systems on a domain $\mathcal{D} \in \mathbb{R}^D$, $D \in \{1, 2, 3\}$, in which the uncertainty sources characterized by Ξ may be represented in one or many relevant parameters, e.g., boundary conditions and/or initial conditions. The solution u is governed by the equations

$$\mathcal{L}(\boldsymbol{x}, t, \Xi; u(\boldsymbol{x}, t, \Xi)) = 0, \quad \boldsymbol{x} \in \mathcal{D},$$
$$\mathcal{I}(\boldsymbol{x}, \Xi; u(\boldsymbol{x}, 0, \Xi)) = 0, \quad \boldsymbol{x} \in \mathcal{D},$$
$$\mathcal{B}(\boldsymbol{x}, t, \Xi; u(\boldsymbol{x}, t, \Xi)) = 0, \quad \boldsymbol{x} \in \partial\mathcal{D}, \tag{10}$$

where $\mathcal{L}$, $\mathcal{I}$, and $\mathcal{B}$ are differential operators depending on the physics of the problem, the initial conditions, and the boundary conditions, respectively. Our objective is to approximate the QoI, here $u(\boldsymbol{x}_0, t_0, \Xi)$, for some fixed spatial location $\boldsymbol{x}_0$ and time t_0. Since we denote the realizations of the random inputs by $\boldsymbol{\xi}^{(i)}$, the corresponding output is $u(\boldsymbol{x}_0, t_0, \boldsymbol{\xi}^{(i)})$. To reduce notation, we drop the reference to $\boldsymbol{x}_0$ and t_0, and simply write $u(\Xi)$ and $u(\boldsymbol{\xi}^{(i)})$.

2.2. Polynomial chaos expansion (PCE)

We rely on the PCE (1) to approximate the QoI, $u(\Xi)$. For convenience, we assume that the input random variables, Ξ_k, are independent and identically distributed according to the probability density function ρ_k, and define $\{\psi_{i_k}(\Xi_k)\}$ to be the complete set of polynomials of degree $i_k \in \mathbb{N} \cup \{0\}$ orthonormal with respect to the weight function ρ_k [34,4]. Hence, the multivariate orthonormal polynomials in Ξ are given by the products of the univariate orthonormal polynomials,

$$\psi_{\boldsymbol{i}}(\Xi) = \prod_{k=1}^{d} \psi_{i_k}(\Xi_k), \tag{11}$$

where each $\boldsymbol{i} \in \{(i_1, \ldots, i_d) : i_k \in \mathbb{N} \cup \{0\}\}$ is a d-dimensional multi-index of non-negative integers. For computation, we truncate the expansion in (1) to the set of P basis functions associated with the subspace of polynomials of total order not greater than p, that is $\sum_{k=1}^{d} i_k \leq p$. For convenience, we also order these P basis functions so that they are indexed by $\{1, \ldots, P\}$, as in (1), where there should be no confusion in using either notation. The basis set $\{\psi_j(\Xi)\}_{j=1}^{P}$ has cardinality

$$P = \frac{(d+p)!}{d!p!}. \tag{12}$$

2.3. ℓ_1-Minimization with gradient information

We generate derivative information for the QoI, denoted by $\boldsymbol{u}_\partial$, and correspondingly evaluate the derivatives of $\psi_j(\Xi)$, $j = 1, \ldots, P$, at the realizations $\boldsymbol{\xi}^{(i)}$, $i = 1, \ldots, N$, stored in a matrix $\boldsymbol{\Psi}_\partial$, as in (7) and (8). For brevity, we define

$$\tilde{\boldsymbol{u}} = \begin{pmatrix} \boldsymbol{u} \\ \boldsymbol{u}_\partial \end{pmatrix} \tag{13}$$

and

$$\tilde{\boldsymbol{\Psi}} = \begin{pmatrix} \boldsymbol{\Psi} \\ \boldsymbol{\Psi}_\partial \end{pmatrix} \boldsymbol{W}, \tag{14}$$

where $\tilde{\boldsymbol{u}} \in \mathbb{R}^{N(d+1)}$, and $\tilde{\boldsymbol{\Psi}} \in \mathbb{R}^{N(d+1) \times P}$ is referred to as the *gradient-enhanced measurement matrix*. Here, $\boldsymbol{W}$ is a positive-diagonal matrix, whose definition is deferred until Section 3.2. Gradient-enhanced ℓ_1-minimization solves the problem

$$\arg\min_{\boldsymbol{c}} \|\boldsymbol{c}\|_1 \quad \text{subject to} \quad \left\| \tilde{\boldsymbol{u}} - \tilde{\boldsymbol{\Psi}} \boldsymbol{c} \right\|_2 \leq \delta, \tag{15}$$

where δ generally differs from the choice in (4) and $\tilde{\boldsymbol{\Psi}}$ is assumed to be normalized such that $\mathbb{E}\left[N^{-1} \tilde{\boldsymbol{\Psi}}^T \tilde{\boldsymbol{\Psi}}\right] = \boldsymbol{I}$, the $P \times P$ identity matrix.

Remark 2.1. As the matrix $\boldsymbol{W}$ renormalizes the PC basis functions used in (1), the coefficient vectors $\boldsymbol{c}$ computed by (4) and (15) differ by these normalizations.

We next begin a theoretical development to justify the gradient-enhanced ℓ_1-minimization approach.

3. Theoretical discussion

We present results supporting the premise that the inclusion of derivative information does not reduce the stability of solutions recovered via (15) when compared to solutions recovered via (4), i.e., in the absence of derivative information. We refer to these solutions and the methods to attain them using the adjectives gradient-enhanced and standard, respectively. We perform our analysis here with Hermite polynomials as they possess the convenient property that they and their derivatives are orthogonal with respect to the same measure, as by (5.5.10) of [35]. While we consider the Probabilists' polynomials here for exposition, the Physicists' polynomials would produce analogous results. Though we do not consider the details here, this analysis may be extended to the case of Laguerre or Jacobi polynomials where the derivative polynomials form an orthogonal system with respect to a measure that differs from the orthogonality measure, yet is still explicitly known, as by (5.1.15) and (4.21.7) of [35], respectively.

First, we motivate and summarize results for standard ℓ_1-minimization. Then we expand on these results in the case of gradient-enhanced ℓ_1-minimization. The path of this analysis flows through the Restricted Isometry Constant (RIC) [5], which is denoted by $\delta_s(\boldsymbol{\Psi})$ and is defined to be the smallest number satisfying

$$(1 - \delta_s(\boldsymbol{\Phi}))\|\boldsymbol{y}\|_2^2 \leq \|\boldsymbol{\Phi}\boldsymbol{y}\|_2^2 \leq (1 + \delta_s(\boldsymbol{\Phi}))\|\boldsymbol{y}\|_2^2, \tag{16}$$

for $\boldsymbol{\Phi} \in \mathbb{R}^{N \times P}$ and all $\boldsymbol{y} \in \mathbb{R}^P$ with at most s non-zero entries. Here, $\delta_s(\boldsymbol{\Phi})$ yields a uniform bound on the spectral radius of the submatrices of $\boldsymbol{\Phi}$ formed by selecting any s columns of $\boldsymbol{\Phi}$. Often, the matrix being considered is clear from context, and we then shorten $\delta_s(\boldsymbol{\Phi})$ to δ_s. Related to the RIC are restricted isometry properties that occur when the RIC reaches a small enough threshold, and guarantee that ℓ_1-minimization with the given matrix is a stable computation. An example of such a restricted isometry property is given in Theorem 3.1 below reported from [39]. This theorem shows that if $\delta_{2s} < 3/(4 + \sqrt{6})$, where s is dictated by the specific problem, then a stable recovery is assured.

Theorem 3.1. (See [39].) Let $\boldsymbol{c} \in \mathbb{R}^P$ represent a solution we seek to approximate, and let $\hat{\boldsymbol{c}}$ be the solution to (4). Let

$$\boldsymbol{\eta} := \boldsymbol{\Psi}\boldsymbol{c} - \boldsymbol{u},$$

denote the contribution from sources of error, and let ϵ from (4) be chosen such that $\|\boldsymbol{\eta}\|_2 < \epsilon$. If

$$\delta_{2s}(\boldsymbol{\Psi}) < \delta_\star := 3/(4 + \sqrt{6}) \approx 0.4652,$$

then the following error estimate holds,

$$\|\boldsymbol{c} - \hat{\boldsymbol{c}}\|_2 \leq \frac{c_1}{\sqrt{s}} \inf_{\|\boldsymbol{c}_s\|_0 \leq s} \|\boldsymbol{c}_s - \boldsymbol{c}\|_1 + c_2 \epsilon,$$

where c_1 and c_2 depend only on δ_{2s}, and $\|\cdot\|_0$ refers to the number of non-zero elements of the vector.

We note that, related to the discussion in Section 3.1, ϵ in Theorem 3.1 may be selected so that $\|\boldsymbol{\eta}\|_2 < \epsilon$ holds with high probability. Unfortunately, identifying the RIC for a given matrix requires a computation for every submatrix of s columns, which is intractable in most situations of interest. As a practical alternative we instead choose to bound the RIC

in a probabilistic sense, allowing us to identify a probability that the RIC is below a chosen threshold. In this way we can guarantee that a restricted isometry property, such as the one of Theorem 3.1, holds with a certain probability. To do so, we introduce a definition of coherence, first considering the standard case, before expanding its definition to the gradient-enhanced case.

3.1. Standard ℓ_1-minimization analysis

We consider here an approach that uses arguments similar to those in [37,9]. We note that those works did not proceed through the RIC as we do here. First, let $\mathcal{Q}$ be an arbitrary subset of the sample space for Ξ, that is the values which Ξ can take. Here, $\mathcal{Q}$ is used to truncate the domain to one on which the basis functions, here Hermite polynomials, can be uniformly bounded. The coherence parameter for the standard approach [37,9] is defined to be

$$\mu_{\mathcal{Q}} := \sup_{k, \boldsymbol{\xi} \in \mathcal{Q}} |\psi_k(\boldsymbol{\xi})|_2^2, \tag{17}$$

which for a precompact $\mathcal{Q}$ is guaranteed to be finite. An example of a $\mathcal{Q}$ suitable for use with Hermite polynomials, and used in [9,16], is

$$\mathcal{Q} := \{\boldsymbol{\xi} : \|\boldsymbol{\xi}\|_2^2 \le (4 + \epsilon_{p,d})p + 2\}, \tag{18}$$

where $\epsilon_{p,d}$ is a positive constant, which may be arbitrarily small but close to zero in an asymptotic analysis of the behavior of Hermite polynomials. We note that this truncation has not been analyzed when the number of samples is exponentially greater than the number of basis functions; however, this is not an issue here where the number of samples is typically less than or not substantially greater than the number of basis functions. The definition of coherence parameter as in (18) leads to the following theorem, taken from Theorem 4.1 of [9].

Theorem 3.2. *For d-dimensional polynomials of order $p \ge 1$, the coherence in (17) is bounded by*

$$\mu_{\mathcal{Q}} \le C_0 \cdot C_1^p, \tag{19}$$

where C_0 and C_1 are modest constants depending on $d, p, \epsilon_{p,d}$. As $p/d \to \infty$, C_0 decreases to 1, and C_1 decreases to a limit of $\exp(2 - \log(2)) \approx 3.7$.

This shows an exponential dependence of the coherence parameter on p. Let $\boldsymbol{X}_k$ denote a row vector consisting only of basis polynomial evaluations at $\boldsymbol{\xi}^{(k)}$. We bound the RIC for the matrix $\boldsymbol{\Psi}$ as defined in (6).

Theorem 3.3. *For any chosen $\mathcal{Q}$, we may bound the RIC in a probabilistic sense by*

$$\mathbb{P}(\delta_s < t) \ge \mathbb{P}(\mathcal{Q})^N - \exp\left(-C_{\mathcal{Q}} \frac{Nt}{s\mu_{\mathcal{Q}}} + s + \log(2s) + s\log(P/s)\right).$$

Remark 3.1. We note that we do not generally require arbitrarily small δ_s, as is seen in Theorem 3.1. The constant $C_{\mathcal{Q}}$ scales with

$$\epsilon_{\mathcal{Q}} := \left\| \mathbb{E}\left(\boldsymbol{X}^T \boldsymbol{X} | \boldsymbol{\xi} \in \mathcal{Q}\right) - \boldsymbol{I} \right\|_2, \tag{20}$$

which is a bias that is negligible in practical contexts for $\mathcal{Q}$ as in (18), [9]. Truncating in this way, neither $C_{\mathcal{Q}}$ nor $\mathbb{P}(\mathcal{Q})$ are problematic in practice.

Remark 3.2. While our primary focus here is ℓ_1-recovery, the RIC corresponding to $s = P$ is useful for analyzing the stability of a least squares solution, and so this result is also applicable to ℓ_2-minimization. For ways in which this parameter may bound error from solutions computed via ℓ_2-minimization we point the interested reader to [38,16]. In this case, a slight adjustment to the proof gives the bound

$$\mathbb{P}(\delta_P < t) \ge \mathbb{P}(\mathcal{Q})^N - 2P \exp\left(-C_{\mathcal{Q}} \frac{Nt}{P\mu_{\mathcal{Q}}}\right).$$

The following corollary, which follows from a rearrangement of the result of Theorem 3.3, highlights the relationship between several quantities.

Corollary 1. *To insure that $\delta_s < \delta_\star$ with probability $p_\star$, it is sufficient to take $N_\star$ satisfying*

$$N_\star \delta_\star \ge \frac{s\mu_{\mathcal{Q}}}{C_{\mathcal{Q}}} \left[s + \log(2s) + s\log(P/s) - \log\left(\mathbb{P}(\mathcal{Q})^{N_\star} - p_\star\right) \right].$$

Remark 3.3. This result has a quadratic scaling in s that is serviceable for many choices of s and N, but is less desirable than a nearly linear scaling in s when considering asymptotic behavior in $N, s,$ and P. Though we leave detailed consideration to the proof of Theorem 3.3, a more careful chaining proof utilizing majorizing measures can reduce this to a linear dependence on s, up to logarithmic factors in s which arise due to structure in the random matrix $\tilde{\Psi}$. For example, following Theorem 8.4 of [39] gives sufficiency if $N_\star$ satisfies

$$\frac{N_\star}{\log(10N_\star)} \geq D_1 \delta_\star^{-2} s \log^2(100s) \mu_{\mathcal{Q}} \log(100P); \tag{21}$$

$$N_\star \geq D_2 \delta^{-2} s \mu_{\mathcal{Q}} \log\left(\left(\mathbb{P}(\mathcal{Q})^{N_\star} - p_\star\right)^{-1}\right). \tag{22}$$

Here $D_1 \leq 50963$, $D_2 \leq 456$, and we note that this bound depends on the squared value of $\delta_\star$. It is appropriate to mention this bound showing a linear dependence on s up to logarithmic factors, for comparison with results in [39–41]. Explicit connection with these works is given in Remark 5.1.

For example, using Corollary 1 and $\delta_\star$ as in Theorem 3.1, gives a guarantee for stability of solutions to (4) when the number of samples N satisfies

$$N \geq \frac{(4 + \sqrt{6})s\mu_{\mathcal{Q}}}{3C_{\mathcal{Q}}}\left[s + \log(2s) + s\log(P/s) - \log\left(\mathbb{P}(\mathcal{Q})^N - p_\star\right)\right].$$

We note that N appears on both the left and right sides of the equation. This arises from the truncation, which has a technical issue when N is exponentially larger than P; specifically, the probability that at least one sample had fallen in $\mathcal{Q}^c$ becomes large, while the analysis relies on this event being rare. As N is very often smaller than P, and rarely chosen to be exponentially larger than P, this issue is not of practical concern. Next, we extend these results to the case where gradient information is included.

3.2. Gradient-enhanced ℓ_1-minimization analysis

In the previous case the rows of Ψ were independent, while $\tilde{\Psi}$ has a more subtle independence structure, as only the sets of rows associated with the independent samples are independent. Related to this point, we let the $(d+1) \times P$ block of independent information related with the kth sample be given by X_k, with a generic realization given by X. Specifically,

$$X(i, j) = \frac{\partial \psi_j}{\partial \Xi_i}(\xi), \quad i = 1, \ldots, d;$$
$$X(d+1, j) = \psi_j(\xi).$$

That is, the last row corresponds to the realizations of the function, while the first d rows correspond to the derivative information. We note that there is no effect in the computations considered here by rearranging the rows of the matrix $\tilde{\Psi}$. The RIC will necessarily be larger if the norms of the matrix columns are different. Adjusting for this can be done by adjusting the basis functions themselves. In the case of standard ℓ_1-minimization, we use orthonormal polynomials. Here, we will wish to include derivatives of those polynomials as well, requiring a different normalization. We use the Probabilists' Hermite polynomials, and the following lemma is used to identify this normalization.

Lemma 3.1. *For d-dimensional orthonormal Probabilists' Hermite polynomials ψ_i and ψ_j with order i_k, j_k in dimension k,*

$$\mathbb{E}\left(\psi_i(\Xi)\psi_j(\Xi) + \sum_{k=1}^d \frac{\partial \psi_i}{\partial \Xi_k}(\Xi)\frac{\partial \psi_j}{\partial \Xi_k}(\Xi)\right) = \delta_{i,j}\left(1 + \sum_{k=1}^d i_k\right), \tag{23}$$

where $\delta_{i,j}$ is the Kronecker Delta.

This suggests a different normalization of the basis functions to enforce that columns of the gradient-enhanced measurement matrix have the same expected ℓ_2-norm. We refer to this normalization as *gradient-normalization*. Specifically, we multiply the orthonormal basis function ψ_i by

$$\tilde{w}_i := \left(1 + \sum_{k=1}^d i_k\right)^{-1/2}, \tag{24}$$

to gradient-normalize. We note that it is these weights that define the W in (9) and (15). In this work, we assume that when derivative information is included, that those basis functions are gradient-normalized, and that when derivative information is not included, that those basis functions are orthonormal, which we refer to as standard-normalization. This insures that the expected norms for the columns of the sampled matrices are consistent in both cases.

Recalling that N denotes the number of samples used, our analysis focuses on the Gramian matrix

$$\boldsymbol{M} := \frac{1}{N} \sum_{k=1}^{N} \boldsymbol{X}_k^T \boldsymbol{X}_k. \tag{25}$$

Here, $\boldsymbol{M} = N^{-1} \boldsymbol{\Psi}^T \boldsymbol{\Psi}$ for the standard approach, while $\boldsymbol{M} = N^{-1} \tilde{\boldsymbol{\Psi}}^T \tilde{\boldsymbol{\Psi}}$ for the gradient-enhanced approach. We now present some summarized results for the standard approach that will be compared to the results presented for the gradient-enhanced case. To analyze the spectrum of $\boldsymbol{M}$ in the case of gradient-enhanced ℓ_1-minimization, we use the following definition, which generalizes the ℓ_1-coherence as studied in [37,9,36] and defined in (17). Let $\mathcal{Q}$ be an arbitrary subset of Ω, which we use to truncate the sample space to insure a uniformly bounded polynomial system, e.g. as in [37,9,16], and let

$$\beta_{\mathcal{Q}} := \sup_{k, \boldsymbol{\xi} \in \mathcal{Q}} \| \boldsymbol{X}(:,k) \|_2^2. \tag{26}$$

This parameter is a generalization of $\mu_{\mathcal{Q}}$ in the case where rows are not independent, but sets of rows are. Note that in the case that $\boldsymbol{X}$ is a row vector, then this definition reduces to (17). Specifically, the results of Section 3.1 all hold when substituted for this parameter, which we highlight as a theorem.

Theorem 3.4. *The theorems of Section 3.1 hold for the gradient-enhanced case when $\mu_{\mathcal{Q}}$ is replaced by $\beta_{\mathcal{Q}}$.*

We conclude our analysis with three results which demonstrate that the inclusion of derivative information, coupled with gradient-normalization, does not reduce stability over the corresponding approach without derivative information.

The first result is an inequality concerning the coherence parameter defined in (17) and (26), showing that including derivative information does not require any weakening of the bounds of Section 3.1 in the case without derivative information. This inequality then directly applies to all of the theorems in Section 3.1.

The second result is a direct null-space comparison of the two different matrices, which is known to be fundamental for recovery of exactly-sparse solutions [5,42,43]. Specifically, as $\| \tilde{\boldsymbol{\Psi}} \boldsymbol{c} - \tilde{\boldsymbol{u}} \| < \delta$ (or $\| \boldsymbol{\Psi} \boldsymbol{c} - \boldsymbol{u} \| < \delta$) is enforced, the difference between potential solutions is close to an element of the null-space of $\tilde{\boldsymbol{\Psi}}$ (or $\boldsymbol{\Psi}$), so reducing the dimension of the null-space correspondingly reduces the space of potential solutions.

The third result concerns a bound on the inner-products of columns of the measurement-matrix. This is related to the RIC in that if the inner-product between several pairs of columns is of large absolute value then a linear combination of those columns will have small norm, resulting in a larger RIC. Similarly, if those inner-products are of small absolute value, then no linear combination will have a small norm. An analogous observation may be made regarding a linear combination of columns having a much larger norm. For this reason it is beneficial if the inner-product between columns is of small absolute value.

Theorem 3.5. *Let $\tilde{\boldsymbol{\Psi}}$ be a realized measurement matrix with derivative information that is gradient-normalized. Similarly, let $\boldsymbol{\Psi}$ be a realized measurement matrix with standard-normalization and no derivative information.*

Assume that $\boldsymbol{\Psi}$ and $\tilde{\boldsymbol{\Psi}}$ are formed from the same realized input samples, $\{\boldsymbol{\xi}^{(i)}\}_{i=1}^N$, so that up to row weighting, $\boldsymbol{\Psi}$ is a sub-matrix of $\tilde{\boldsymbol{\Psi}}$. Then the following statements related to the recovery of solutions via ℓ_1-minimization hold.

R1. *Using the definition in (26) for the two different approaches,*

$$\beta_{\mathcal{Q}}(\tilde{\boldsymbol{\Psi}}) \le \mu_{\mathcal{Q}}(\boldsymbol{\Psi}),$$

and this inequality is almost-surely strict.

R2. *If $\mathcal{N}(\cdot)$ represents the null-space, then $\mathcal{N}(\tilde{\boldsymbol{\Psi}} \boldsymbol{W}^{-1}) \subset \mathcal{N}(\boldsymbol{\Psi})$, and this is almost-surely a strict subset when $\boldsymbol{\Psi}$ is undersampled. Specifically, it almost-surely holds that $\dim(\mathcal{N}(\tilde{\boldsymbol{\Psi}})) = \max\{0, P - (d+1)N\}$ while $\dim(\mathcal{N}(\boldsymbol{\Psi})) = \max\{0, P - N\}$.*

R3. *If subscripts of matrices correspond to the columns, and i_k denotes the order of the basis polynomial in the $\boldsymbol{i}$th column in the kth dimension, then the associated inner product of columns is bounded by,*

$$\sup_{\boldsymbol{i} \ne \boldsymbol{j}} |(\tilde{\boldsymbol{\Psi}}_{\boldsymbol{i}}, \tilde{\boldsymbol{\Psi}}_{\boldsymbol{j}})| \le \sup_{\boldsymbol{i} \ne \boldsymbol{j}} \frac{|(\boldsymbol{\Psi}_{\boldsymbol{i}}, \boldsymbol{\Psi}_{\boldsymbol{j}})|(1 + \sum_{k=1}^{d} \sqrt{i_k j_k})}{\sqrt{(1 + \sum_{k=1}^{d} i_k)(1 + \sum_{k=1}^{d} j_k)}} \le \sup_{\boldsymbol{i} \ne \boldsymbol{j}} |(\boldsymbol{\Psi}_{\boldsymbol{i}}, \boldsymbol{\Psi}_{\boldsymbol{j}})|. \tag{27}$$

Remark 3.4. The theorem is presented for full gradient information to ease presentation, but the three points generalize to the case that derivative information is included for a fraction of samples. Specifically, (1) holds with an appropriate adjustment to the dimensionality, (2) holds with an adjustment to the basis dependent multiplicative constant, and (3) holds if $\tilde{\boldsymbol{\Psi}}$ has derivative information for only a few samples. In summary adding derivative information for even a percentage of the samples leads to bounds as in Theorem 3.5.

3.3. A note on potentially contrasting solutions

It is of practical importance to note what functions are recovered in an asymptotic sense by the gradient-enhanced and standard approaches, as they may differ significantly. The gradient-enhanced method gives $\hat{u}$ approximating u in a Sobolev type loss function,

$$L(\hat{u}, u) := \|\hat{u} - u\|_{\ell_2(\Xi, N)}^2 + \sum_{k=1}^{d} \left\| \frac{\partial(\hat{u} - u)}{\partial \Xi_k} \right\|_{\ell_2(\Xi, N)}^2, \tag{28}$$

where the $\ell_2(\Xi, N)$ indicates the discrete ℓ_2 norm using N evaluations drawn from realizations of Ξ, here $(\xi^{(1)}, \ldots, \xi^{(N)})$. This norm is also normalized by N^{-1} so that as N goes to infinity this norm tends to $\mathcal{L}_2(\Xi)$, the standard $\mathcal{L}_2$ norm associated with the distribution of Ξ. Specifically (15) guarantees that $NL(u, \hat{u}) < \delta$. In contrast, without derivative information using standard-normalization and producing a solution via (4), gives $\hat{u}$ approximating u such that $N\|\hat{u} - u\|_{\ell_2(\Xi, N)}^2 < \delta$, that is the partial derivatives of the approximation need not be approximated by those of the target function. As these loss functions differ, so too does the limiting solution produced by each method for a finite expansion order p. However, the sequence of approximations given by the two loss functions and by using increasing p (and accordingly N) will converge to u in the $\mathcal{L}_2(\Xi)$ sense as $p, N \to \infty$.

4. Numerical results

In this section, we empirically demonstrate the gradient-enhanced ℓ_1-minimization approach using Hermite polynomials and Gaussian inputs. This is done via three numerical examples: a manufactured PCE; an elliptic PDE with stochastic coefficient; and a plane Poiseuille flow with random boundaries. To compare the standard and gradient-enhanced ℓ_1-minimization solutions, we define an equivalent sample size $\tilde{N}$,

$$\tilde{N} := N_e + \nu N_g, \tag{29}$$

which accounts for the added cost of computing the derivative information. In (29), N_e is the number of samples without derivative information, N_g is the number of samples with derivatives (along all d directions), and ν is a positive parameter depending on the problem at hand and the approach employed to compute the derivatives. For the example of Section 4.2, the cost of generating d derivatives of the QoI, obtained by the adjoint sensitivity method, is roughly the same as that of evaluating the QoI, thus implying that $\nu = 2$. For transient problems where the cost of solving the adjoint equations for derivative calculations may be considerably more than that of a single QoI evaluation, $\nu > 2$. Nevertheless, we here present all cost comparisons in terms of the number of equivalent sample size $\tilde{N}$ in (29), for choices of ν that we shall specify. For simulations based on standard ℓ_1-minimization, we set $\tilde{N} = N_e$. Additionally, for the interest of convenience, we ignore the cost of solving the ℓ_1-minimization problems in (4) and (15). This is a valid assumption as often that cost is negligible relative to the cost of evaluating the QoI or its derivatives. For terminology, if $X\%$ of samples used in the computation of the solution contain derivative information, then we say that method is $X\%$ gradient-enhanced. In this way, the standard approach without derivatives would give a solution that is 0% gradient-enhanced, denoted as standard in the subsequent figures. Similarly, if all samples include derivative information, the associated solution would be 100% gradient-enhanced. In all cases, all samples are drawn from the orthogonality distribution, that is the multivariate standard normal distribution.

4.1. Case I: A manufactured PCE

First, we consider the reconstruction of a manufactured PCE, in which the sparsity and the entries of the coefficient vector c are *a priori* prescribed. Specifically, we set the dimension of the expansion to $d = 25$ and use a $p = 3$ order PCE (hence $P = 3276$ basis functions) to manufacture the QoI $u(\Xi)$. To generate c, we first draw its P entries independently from the standard Gaussian distribution. We then retain $|\mathcal{C}| \in \{50, 150\}$ coefficients with largest magnitude and set the rest to zero. This gives a randomized sparsity support. Finally, the realizations of $u(\Xi)$ and its derivative with respect to Ξ_k, $k = 1, \ldots, d$, are generated by
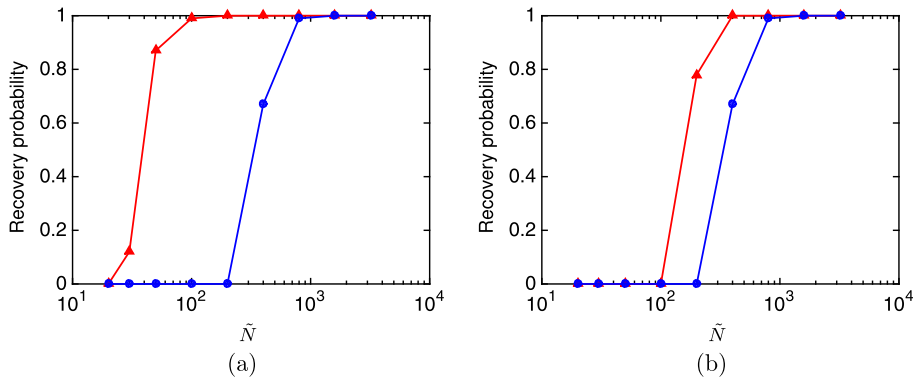
$$u(\xi^{(i)}) = \sum_{j=1}^{P} c_j \psi_j(\xi^{(i)}) \tag{30}$$
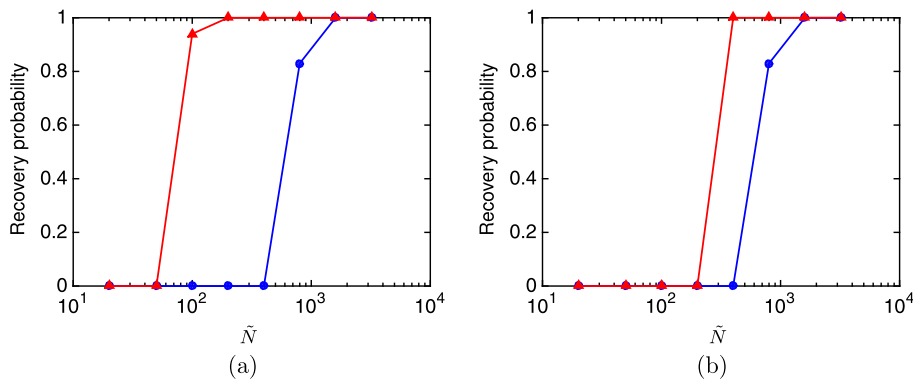
and

$$u_{\partial_k}(\xi^{(i)}) := \sum_{j=1}^{P} c_j \frac{\partial \psi_j}{\partial \Xi_k}(\xi^{(i)}), \tag{31}$$

respectively. We then approximate the PCE coefficients c via (4) and (15) from these generated data.

Fig. 1. Probability of successful recovery of manufactured PCE with $p = 3$, $d = 25$, and sparsity $|\mathcal{C}| = 50$ via gradient-enhanced and standard ℓ_1-minimization. (a) 100% vs. 0% gradient-enhanced. (b) 20% vs. 0% gradient-enhanced. ($-\blacktriangle-$ gradient-enhanced, $-\bullet-$ standard.)

Fig. 2. Probability of successful recovery of manufactured PCE with $p = 3$, $d = 25$, and sparsity $|\mathcal{C}| = 150$ via gradient-enhanced and standard ℓ_1-minimization. (a) 100% vs. 0% gradient-enhanced. (b) 20% vs. 0% gradient-enhanced. ($-\blacktriangle-$ gradient-enhanced, $-\bullet-$ standard.)

4.1.1. Results

Beginning with the $|\mathcal{C}| = 50$ case, we seek to recover the manufactured $\boldsymbol{c}$. If the computed solution, denoted by $\hat{\boldsymbol{c}}$, has a relative root-mean-square-error (RRMSE) below 0.01%, then we call it a successful recovery of $\boldsymbol{c}$.
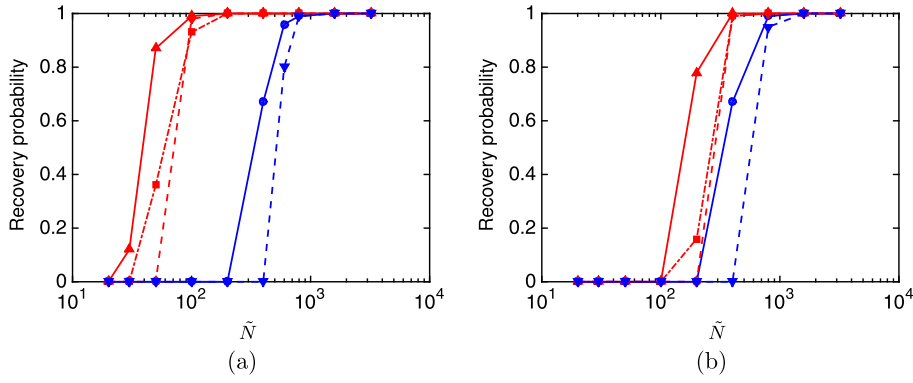
We first consider the case where evaluations of $u(\boldsymbol{\Xi})$ and its derivatives are exact, i.e., noise free. In Fig. 1, we compare the probability of successful recovery for gradient-enhanced and standard approaches, using 100 independent replications for each $\tilde{N}$. To set $\tilde{N}$, we pretend the cost of evaluating d derivatives of $u(\boldsymbol{\Xi})$ is the same as that of evaluating $u(\boldsymbol{\Xi})$, and therefore we set $\nu = 2$.

In Fig. 1(a), we see that 100% gradient-enhanced ℓ_1-minimization helps in reducing the computational effort to recover $\boldsymbol{c}$, while Fig. 1(b) demonstrates a notable, but less, improvement for 20% gradient-enhanced ℓ_1-minimization. The figure suggests that adding derivative information is a cost-effective means to increase the probability of successfully recovering $\boldsymbol{c}$ at low sample sizes $\tilde{N}$.
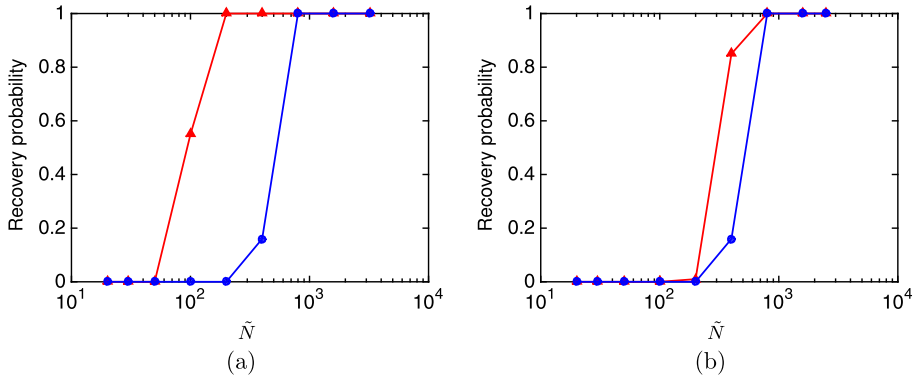
Fig. 2 shows similar results for sparsity $|\mathcal{C}| = 150$. Since $|\mathcal{C}|$ is larger in this case, both standard and gradient-enhanced methods require more samples for recovery with a given success probability. This is consistent with the sampling rate presented in Corollary 1. We, however, note that the inclusion of derivative information still enhances the solution recovery.

In practice, there is often error (or noise) in the evaluation of $u(\boldsymbol{\Xi})$ and its derivatives, with the latter being more prone to errors. To model such inaccuracies, here we multiply the realizations of $u(\boldsymbol{\Xi})$ and their derivatives from (30) and (31), respectively, by independent realizations of $(1 + \epsilon_N)$, where ϵ_N is a zero mean Gaussian random variable with variance 10^{-5}. In Fig. 3, we consider the sparsity $|\mathcal{C}| = 50$ case and compute the probability of successful solution recovery as a function of $\tilde{N}$. Similar as in the previous test cases, we use an RRMSE error of 0.01% to identify a successful recovery. We observe that the inclusion of derivative information, while imprecise, still improves the performance of the standard ℓ_1-minimization.

Fig. 4 considers a case with lower dimension $d = 8$, and higher order $p = 6$, giving $P = 3003$. This shows the effect with regards to having lower dimensionality and higher order on recovery. It is noted that increasing the order of Hermite expansions significantly increases both $\mu_{\mathcal{Q}}$ in (17) and $\beta_{\mathcal{Q}}$ in (26). While tangible improvement in solution recovery is still noticed when considering gradient-enhanced samples, the amount of improvement is reduced. This reduction in efficacy may be related to the reduction in the dimensionality of information provided by the gradient.

Fig. 3. Probability of successful recovery of gradient-enhanced and standard ℓ_1-minimization for the manufactured PCE case with $p = 3$, $d = 25$, and sparsity $|\mathcal{C}| = 50$. (a) 100% gradient-enhanced. (b) 20% gradient-enhanced. (—▲— gradient-enhanced, --■-- gradient-enhanced with noisy $\boldsymbol{u}_\partial$ only, --◆-- gradient-enhanced with both noisy $\boldsymbol{u}$ and $\boldsymbol{u}_\partial$, —●— standard, --▼-- standard with noisy $\boldsymbol{u}$.)

Fig. 4. Probability of successful recovery of gradient-enhanced and standard ℓ_1-minimization for the manufactured PCE case with $p = 8$, $d = 6$, and sparsity $|\mathcal{C}| = 50$. (a) 100% vs. 0% gradient-enhanced. (b) 20% vs. 0% gradient-enhanced. (—▲— gradient-enhanced, —●— standard.)

4.2. Case II: Two-dimensional elliptic PDE with random coefficient

We next consider the two-dimensional (in space) elliptic PDE

$$-\nabla \cdot (a(\boldsymbol{x}, \boldsymbol{\Xi})\nabla u(\boldsymbol{x}, \boldsymbol{\Xi})) = 1, \quad \boldsymbol{x} \in \mathcal{D} = [0, 1]^2,$$
$$u(\boldsymbol{x}, \boldsymbol{\Xi}) = 0, \quad \boldsymbol{x} \in \partial\mathcal{D}, \tag{32}$$

where the diffusion coefficient $a(\boldsymbol{x}, \boldsymbol{\Xi})$ is modeled by the lognormal random field

$$a(\boldsymbol{x}, \boldsymbol{\Xi}) = \exp\left[\bar{a} + \sigma_a \sum_{k=1}^{d} \sqrt{\lambda_k}\phi_k(\boldsymbol{x})\Xi_k\right]. \tag{33}$$

Here, $d = 30$, and Ξ_k, $k = 1, \ldots, d$, are independent standard Gaussian random variables. In addition, $\bar{a} = 0.1$, $\sigma_a = 0.5$, and $\{\phi_k\}_{k=1}^{d}$ are the eigenfunctions corresponding to the d largest eigenvalues $\{\lambda_k\}_{k=1}^{d}$ of the Gaussian covariance kernel
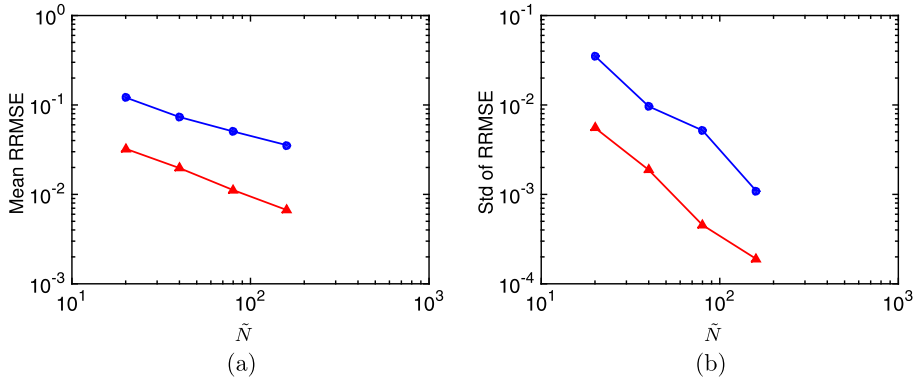
$$C_{aa}(\boldsymbol{x}_1, \boldsymbol{x}_2) = \exp\left[-\frac{(x_1 - x_2)^2}{l_c^2} - \frac{(y_1 - y_2)^2}{l_c^2}\right], \tag{34}$$

with correlation length $l_c = 1/16$.

The QoI $u(\boldsymbol{\Xi})$ is chosen as the solution to (32) at location $\boldsymbol{x} = (0.5, 0.5)$, i.e., $u(\boldsymbol{\Xi}) = u((0.5, 0.5), \boldsymbol{\Xi})$. The derivatives $\partial u/\partial \Xi_k$ are computed using the adjoint sensitivity method explained in detail in Section 4.2.1. Both forward and adjoint solvers are implemented in the finite element method (FEM) project FEniCS [44].

4.2.1. Adjoint sensitivity derivatives

The adjoint sensitivity methods are commonly used in research areas such as sensitivity analysis [45,46], optimization [47,48], shape design [49,50], etc., to compute the derivatives of the solutions of interest with respect to the underlying

Fig. 5. Comparison of the statistics of the RRMSE in reconstructing $u\,((0.5, 0.5), \Xi)$, where the realizations of u and its derivatives are computed on a uniform 256×256 FEM mesh. (a) Mean of RRMSE. (b) Standard deviation of RRMSE. (—▲— 100% gradient-enhanced, —●— standard.)

model parameters. In this work, we adopt the discrete adjoint sensitivity method to compute derivatives of the QoI at the Ξ samples.

In detail, for the interest of convenience, we consider the discrete formulation of the generic PDE in (10), at a fixed time, given by the residual equation

$$\mathcal{R}(\boldsymbol{w}, \Xi) = \mathbf{0}, \tag{35}$$

where $\boldsymbol{w} \in \mathbb{R}^M$ contains the discrete values of solution over the spatial domain $\mathcal{D}$, and M is the number of solution degrees-of-freedom. Recalling that $u(\Xi)$ (here a scalar functional of the solution) denotes the QoI, we seek to compute the sensitivity derivatives $du/d\Xi_k$ from

$$\frac{du}{d\Xi_k} = \frac{\partial u}{\partial \Xi_k} + \frac{\partial u}{\partial \boldsymbol{w}} \frac{\partial \boldsymbol{w}}{\partial \Xi_k}. \tag{36}$$

Taking the derivative of (35) with respect to Ξ_k, we have

$$\frac{\partial \mathcal{R}}{\partial \Xi_k} + \frac{\partial \mathcal{R}}{\partial \boldsymbol{w}} \frac{\partial \boldsymbol{w}}{\partial \Xi_k} = \mathbf{0}, \tag{37}$$

which results in $\partial \boldsymbol{w}/\partial \Xi_k = -(\partial \mathcal{R}/\partial \boldsymbol{w})^{-1} \partial \mathcal{R}/\partial \Xi_k$. Plugging this in (36) gives

$$\frac{du}{d\Xi_k} = \frac{\partial u}{\partial \Xi_k} - \frac{\partial u}{\partial \boldsymbol{w}} \left(\frac{\partial \mathcal{R}}{\partial \boldsymbol{w}} \right)^{-1} \frac{\partial \mathcal{R}}{\partial \Xi_k},$$

which can be rewritten as

$$\frac{du}{d\Xi_k} = \frac{\partial u}{\partial \Xi_k} + \boldsymbol{\lambda}^T \frac{\partial \mathcal{R}}{\partial \Xi_k}, \tag{38}$$

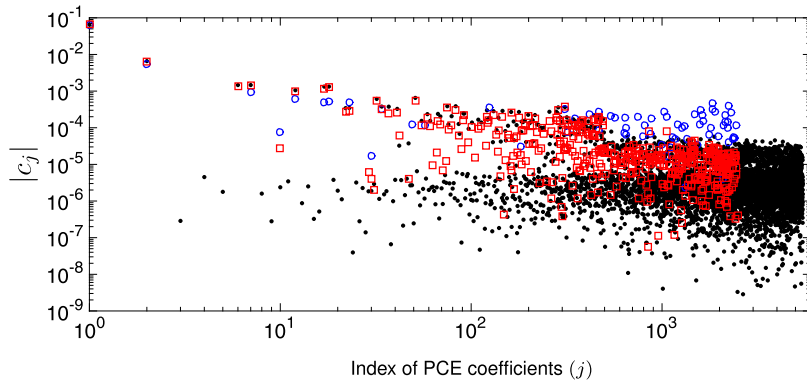where $\boldsymbol{\lambda}$ is the solution to the discrete adjoint equation

$$\left(\frac{\partial \mathcal{R}}{\partial \boldsymbol{w}} \right)^T \boldsymbol{\lambda} = -\left(\frac{\partial u}{\partial \boldsymbol{w}} \right)^T. \tag{39}$$

For the case of elliptic PDE (32), $\partial \mathcal{R}/\partial \boldsymbol{w}$ is the symmetric stiffness matrix of the FEM discretization and $\partial \mathcal{R}/\partial \Xi_k$ in (38) can be computed semi-analytically from (33) and the FEM formulation of (32). Here, we assume the inverse or a factorization of $\partial \mathcal{R}/\partial \boldsymbol{w}$ is not stored when solving for $\boldsymbol{w}$, and that the total cost of obtaining $\partial \mathcal{R}/\partial \Xi_k$ is smaller than that of computing $\boldsymbol{w}$. Therefore, the cost of solving for $\boldsymbol{\lambda}$ from (39) is roughly the same as solving for $\boldsymbol{w}$, which in turn suggests that $\nu = 2$ in (29).

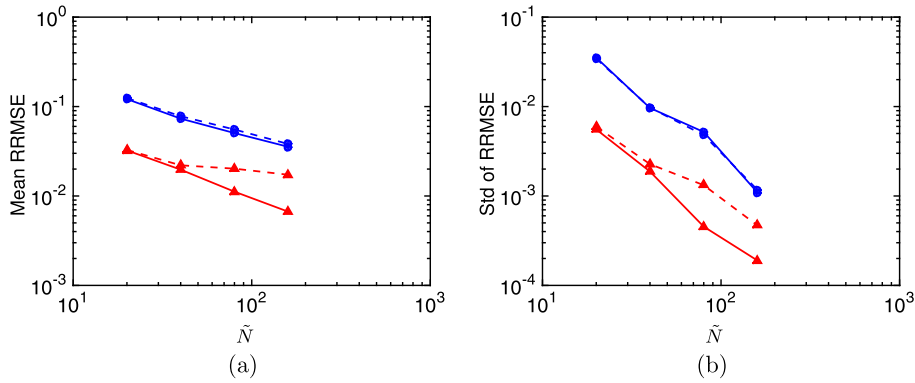### 4.2.2. Results

We approximate $u(\Xi)$ in a Hermite PCE with total degree $p = 3$, and seek to approximate the first 2500 coefficients. We sort the elements of $\{\psi_j\}$ such that, for any given total order basis, the random variables Ξ_k with smaller indices k contribute first to the basis.
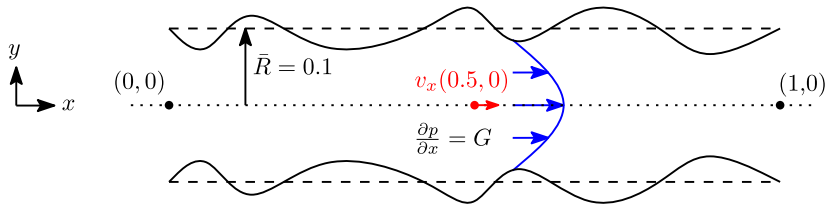
To solve for the coefficients, we first generate the realizations $\boldsymbol{u}$ and $\boldsymbol{u}_\partial$ using a 256×256 uniform, linear FEM mesh, which resolves both quantities with *low* numerical errors. In Fig. 5, we compare the mean and standard deviation of the RRMSE for solutions computed by the standard and gradient-enhanced ℓ_1-minimization, using 100 independent replications. The reference PCE coefficients are computed using least squares regression [16] with 10 000 solution realizations and yields a relative error of 0.36% for 1000 additional validation samples. From Fig. 5, we observe that higher accuracies are achieved

Fig. 6. Approximate PCE coefficients of $u((0.5, 0.5), \Xi)$ with $\tilde{N} = 80$ vs. the reference coefficients obtained by least squares regression. (● reference, ○ standard ℓ_1-minimization, □ gradient-enhanced ℓ_1-minimization.)

Fig. 7. Comparison of the statistics of the RRMSE in reconstructing $u((0.5, 0.5), \Xi)$ via gradient-enhanced and standard ℓ_1-minimization. (a) Mean of RRMSE. (b) Standard deviation of RRMSE. (—▲— 100% gradient-enhanced, —●— standard. Dashed and solid lines, respectively, correspond to the 16×16 and 256×256 mesh simulations.)

Fig. 8. Schematic figure of plane Poiseuille flow with random boundaries.

by the gradient-enhanced ℓ_1-minimization with the same number of samples $\tilde{N}$. In Fig. 6, we show the magnitude of the approximate PCE coefficients computed via standard and gradient-enhanced ℓ_1-minimization with $\tilde{N} = 80$ samples. More accurate coefficient estimates are obtained by the gradient-enhanced ℓ_1-minimization.

To study how the accuracy of derivative information affects the accuracy of solution obtained by the gradient-enhanced ℓ_1-minimization, we repeat this experiment on a coarser 16×16 mesh, where the derivative information is noticeably less accurate. We observe from Fig. 7 that, the accuracy improvement achieved by the gradient-enhanced ℓ_1-minimization is not as considerable as in the case of 256×256 mesh. This suggests that high accuracy on the derivative samples $\boldsymbol{u}_\partial$ may be required for the gradient-enhanced ℓ_1-minimization to be most effective.

4.3. Case III: Plane Poiseuille flow with random boundaries

We consider a 2-D plane Poiseuille flow with random boundaries as depicted in Fig. 8.

The average width of the channel is $2\bar{R} = 0.2$, and the width of the channel is model as $2R(x) = 2\bar{R} + 2r(x, \Xi)$, where $2r(x, \Xi)$ describes the random fluctuation of the channel width around $2\bar{R}$ (see Fig. 9). We use $d = 20$ independent standard Gaussian random variables $\Xi_k, k = 1, \ldots, d$, to represent the uncertainty in R, and let
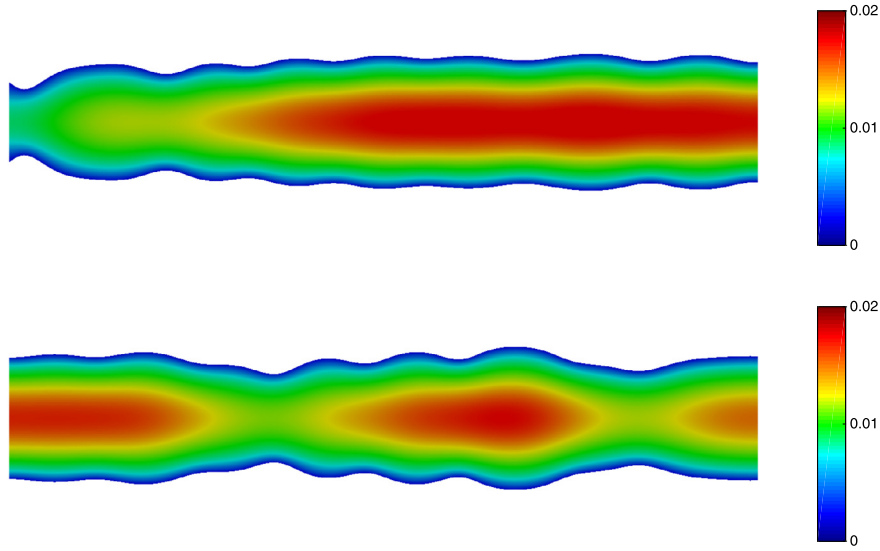
Fig. 9. Two realizations of velocity magnitude corresponding to two realizations of random boundaries.

$$r(x, \mathbf{\Xi}) = \exp\left[\bar{r} + \sigma_r \sum_{k=1}^{d} \sqrt{\lambda_k}\phi_k(x)\,\Xi_k\right]. \tag{40}$$

Here, $\bar{r} = -4$, $\sigma_r = 0.5$, and $\{\lambda_k\}_{k=1}^{d}$ and $\{\phi_k\}_{k=1}^{d}$ are the d largest eigenvalues and corresponding eigenfunctions of the exponential covariance kernel

$$C_{rr}(x_1, x_2) = \exp\left(-\frac{|x_1 - x_2|}{l_c}\right), \tag{41}$$

where the correlation length $l_c = 1/21$.

We seek to investigate the steady state velocity field of the flow, which is governed by the incompressible steady state Navier–Stokes equations

$$
\begin{aligned}
(\mathbf{v} \cdot \nabla)\mathbf{v} - \frac{1}{Re}\nabla^2 \mathbf{v} &= -\nabla p, \quad && \mathbf{x} \in \mathcal{D}(\mathbf{\Xi}), \\
\nabla \cdot \mathbf{v} &= 0, && \mathbf{x} \in \mathcal{D}(\mathbf{\Xi}), \\
\frac{\partial p}{\partial x} &= G, && \mathbf{x} \in \mathcal{D}(\mathbf{\Xi}), \\
\mathbf{v} &= 0, && y = \pm R,
\end{aligned}
\tag{42}
$$

where the Reynolds number $Re = 60$, and p denotes pressure. Notice that in (42), we assume that all physical quantities are non-dimensional. The flow is driven by a pressure gradient $G = -0.1$.

The QoI is $v_x(0.5, 0)$, the horizontal velocity at $(0.5, 0)$, and we approximate it in a Hermite PCE of total degree $p = 3$.
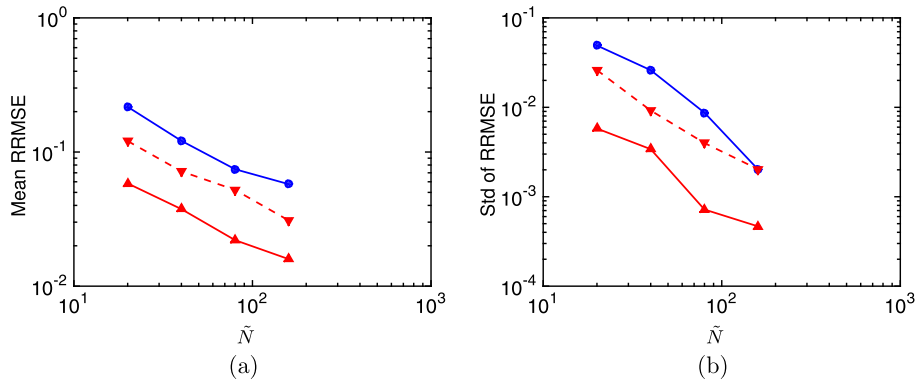
4.3.1. Adjoint sensitivity derivatives

To compute the derivative information, we again adopt the adjoint sensitivity method, in which we approximate $\partial \mathcal{R}/\partial \Xi_k$, $k = 1, \ldots, d$, in (38) via finite difference quotient

$$\frac{\partial \mathcal{R}}{\partial \Xi_k} \approx \frac{\mathcal{R}(\mathbf{w}, \mathbf{\Xi} + \Delta\mathbf{\Xi}^k) - \mathcal{R}(\mathbf{w}, \mathbf{\Xi})}{\varepsilon}. \tag{43}$$

Here, the mth entry of the vector $\Delta\mathbf{\Xi}^k$ is defined as

$$\Delta\Xi_m^k = \begin{cases} \varepsilon, & m = k, \\ 0, & m \neq k. \end{cases} \tag{44}$$

To solve the non-linear problem (42), we employ a standard Newton solver, in which $\partial \mathcal{R}/\partial \mathbf{w}$ from (39) is the Jacobian matrix. In (43), the discrete representation of the velocity field, $\mathbf{w}$, is kept unchanged; hence, (42) does not need to be solved again. The perturbed residual $\mathcal{R}(\mathbf{w}, \mathbf{\Xi} + \Delta\mathbf{\Xi}^k)$ is computed by deforming the mesh to conform to the geometry corresponding to $\mathbf{\Xi} + \Delta\mathbf{\Xi}^k$, without recomputing $\mathbf{w}$. Compared to solving the adjoint equation (39), the cost of calculating

Fig. 10. Comparison of the statistics of the RRMSE in reconstructing $v_x((0.5, 0), \boldsymbol{\Xi})$ via standard and gradient-enhanced ℓ_1-minimization, with 100 independent replications. (a) Mean of RRMSE. (b) Standard deviation of RRMSE. (——▲—— 100% gradient-enhanced, - -▼- - 20% gradient-enhanced, ——●—— standard.)

(43) is negligible. Additionally, in our deterministic solver, computing $\boldsymbol{w}$ required on average 3 Newton steps. Therefore, the extra cost of computing derivative information is roughly equivalent to 1/3 of the cost of computing $\boldsymbol{w}$, which in turn suggests setting $\nu = 4/3$ in (29).

4.3.2. Results

We consider 0%, 20% and 100% gradient-enhanced ℓ_1-minimization, with $\tilde{N} \in \{20, 40, 80, 160\}$.

Fig. 10 displays the comparisons of the mean and standard deviation of the RRMSE, with 100 independent replications of computed PCE coefficients, showing that gradient-enhanced ℓ_1-minimization again leads to cost-effective accuracy improvement.

5. Proofs

5.1. Theorem 3.4

We prove our results here in the case of $\beta_{\mathcal{Q}}$ defined as in (26), which due to the non-independent rows is a slight generalization of the results using (17). Adjusting the proofs present here to account for the latter case requires only the substitution of $\mu_{\mathcal{Q}}$ for $\beta_{\mathcal{Q}}$, showing this Theorem.

5.2. Lemma 3.1

We begin by providing a brief proof for Lemma 3.1, which follows directly from the explicit form for the Hermite derivative as in (5.5.10) of [35].

Proof. Note that linearity of expectation allows us to take the expectation inside the sum, so that we may work on each term independently. For one-dimensional orthonormal Hermite polynomials, where i represents the order of the polynomial, (5.5.10) of [35] shows that

$$\frac{\partial \psi_i}{\partial \Xi}(\Xi) = \sqrt{i}\psi_{i-1}(\Xi).$$ (45)

As the tensor product of orthonormal polynomials is orthonormal, (23) for $\boldsymbol{i} = \boldsymbol{j}$ follows from the derivative being exactly in the direction of a Hermite polynomial. For $\boldsymbol{i} \neq \boldsymbol{j}$, (23) follows because each term in the sum is the expectation of a product of orthogonal polynomials that differ in at least one coordinate, and hence the integral is equal to zero. □

5.3. Theorem 3.3

To prove Theorem 3.3, we appeal to a matrix variant of the Chernoff bound [51], which is similar to approaches taken in [38,37,52,16].

Proof. Recall that $\mathcal{Q}$ represents a truncation for the domain of $\boldsymbol{\Xi}$, and may be given by (18). We have that

$$\boldsymbol{I} = \mathbb{E}\left(\boldsymbol{X}^T\boldsymbol{X}|\boldsymbol{\xi} \in \mathcal{Q}\right)\mathbb{P}(\mathcal{Q}) + \mathbb{E}\left(\boldsymbol{X}^T\boldsymbol{X}|\boldsymbol{\xi} \in \mathcal{Q}^c\right)\mathbb{P}(\mathcal{Q}^c).$$ (46)

A brief calculation gives that

$$\epsilon_{\mathcal{Q}} := \left\| \mathbb{E}\left(\boldsymbol{X}^T \boldsymbol{X} | \boldsymbol{\xi} \in \mathcal{Q} \right) - \boldsymbol{I} \right\|_2 \tag{47}$$

$$\leq \frac{\mathbb{P}(\mathcal{Q}^c)}{\mathbb{P}(\mathcal{Q})} \left(\left\| \mathbb{E}\left(\boldsymbol{X}^T \boldsymbol{X} \Big| \boldsymbol{\xi} \in \mathcal{Q}^c \right) \right\|_2 + 1 \right) \tag{48}$$

bounds the bias introduced from not accepting samples within $\mathcal{Q}^c$. We note that with the truncation in (18), and using this bound, $\epsilon_{\mathcal{Q}} \leq 0.1/\sqrt{P}$ [9], for all N considered here, and in most similar problems. Specifically, an analytic issue does not arise until consideration of exponential levels of oversampling. This analytic issue arises as the truncated events are no longer negligibly rare when a high number of samples are taken.

Restating (47), let $\lambda_{\min}(\cdot)$ and $\lambda_{\max}(\cdot)$ correspond to the smallest and largest eigenvalues of the argument matrix, respectively. Then for any arbitrary set of columns, denoted $\mathcal{S}$,

$$1 - \epsilon_{\mathcal{Q}} \leq \lambda_{\min}\left(\mathbb{E}\left(\boldsymbol{X}^T(:, \mathcal{S}) \boldsymbol{X}(:, \mathcal{S}) | \boldsymbol{\xi} \in \mathcal{Q} \right) \right) \tag{49}$$

$$\leq \lambda_{\max}\left(\mathbb{E}\left(\boldsymbol{X}^T(:, \mathcal{S}) \boldsymbol{X}(:, \mathcal{S}) | \boldsymbol{\xi} \in \mathcal{Q} \right) \right) \leq 1 + \epsilon_{\mathcal{Q}}. \tag{50}$$

From (26) we have that if each sample $\boldsymbol{\xi}^{(k)} \in \mathcal{Q}$, then for all k,

$$\|\boldsymbol{X}_k(:, \mathcal{S})\|_2^2 \leq s\beta_{\mathcal{Q}}, \tag{51}$$

holds uniformly for all choices of $\mathcal{S}$ such that $|\mathcal{S}| < s$. This provides an upper bound on the singular values of our independent self-adjoint matrices, $\boldsymbol{X}_k^T(:, \mathcal{S})\boldsymbol{X}_k(:, \mathcal{S})$, uniformly over all choices of $\mathcal{S}$ consisting of at most s elements. We define

$$\boldsymbol{M}_{\mathcal{S}} := \frac{1}{N} \sum_{k=1}^{N} \boldsymbol{X}_k^T(:, \mathcal{S}) \boldsymbol{X}_k(:, \mathcal{S}). \tag{52}$$

An application of the Chernoff bound as in Theorem 1.1 of [51] and Theorem 1 of [38] gives that for $\delta \in [0, 1]$ and $|\mathcal{S}| \leq s$,

$$\mathbb{P}\left(\lambda_{\min}(\boldsymbol{M}_{\mathcal{S}}) \leq (1 - \delta)(1 - \epsilon_{\mathcal{Q}}) \Big| \boldsymbol{\xi}^{(k)} \in \mathcal{Q} \, \forall k \right) \leq s \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^{\frac{N(1-\epsilon_{\mathcal{Q}})}{s\beta_{\mathcal{Q}}}}; \tag{53}$$

$$\mathbb{P}\left(\lambda_{\max}(\boldsymbol{M}_{\mathcal{S}}) \geq (1 + \delta)(1 + \epsilon_{\mathcal{Q}}) \Big| \boldsymbol{\xi}^{(k)} \in \mathcal{Q} \, \forall k \right) \leq s \left(\frac{e^{\delta}}{(1 + \delta)^{1+\delta}} \right)^{\frac{N(1+\epsilon_{\mathcal{Q}})}{s\beta_{\mathcal{Q}}}}. \tag{54}$$

Note that

$$(1 - \delta)(1 - \epsilon_{\mathcal{Q}}) \geq 1 - t \Rightarrow \delta \leq \frac{t - \epsilon_{\mathcal{Q}}}{1 - \epsilon_{\mathcal{Q}}}; \tag{55}$$

$$(1 + \delta)(1 + \epsilon_{\mathcal{Q}}) \leq 1 + t \Rightarrow \delta \leq \frac{t - \epsilon_{\mathcal{Q}}}{1 + \epsilon_{\mathcal{Q}}}, \tag{56}$$

and so we have a critical δ, given by

$$\delta_t := (t - \epsilon_{\mathcal{Q}})/(1 + \epsilon_{\mathcal{Q}}), \tag{57}$$

is such that for all $\delta < \delta_t$ the matrix $\boldsymbol{M}_{\mathcal{S}}$ is guaranteed to satisfy $\|\boldsymbol{M}_{\mathcal{S}} - \boldsymbol{I}\|_2 \leq t$. Note that for $0 \leq \delta < 1$,

$$\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \geq \frac{e^{\delta}}{(1 + \delta)^{1+\delta}}, \tag{58}$$

and so we may bound the sum of the probabilities by

$$\mathbb{P}\left(\|\boldsymbol{M}_{\mathcal{S}} - \boldsymbol{I}\| \leq t \Big| \boldsymbol{\xi}^{(k)} \in \mathcal{Q} \, \forall k \right) \leq 2s \left(\frac{e^{-\delta_t}}{(1 - \delta_t)^{1-\delta_t}} \right)^{\frac{N(1-\epsilon_{\mathcal{Q}})}{s\beta_{\mathcal{Q}}}}. \tag{59}$$

We now bound this probability. We note that if $\epsilon_{\mathcal{Q}} \leq t$ then

$$0 < \delta_t \leq \frac{t - \epsilon_{\mathcal{Q}}}{1 + \epsilon_{\mathcal{Q}}} \leq t - \epsilon_{\mathcal{Q}}. \tag{60}$$

To create a bound without explicit dependence on $\epsilon_{\mathcal{Q}}$, we note that for

$$c_t := t - \epsilon_{\mathcal{Q}} + (t + \epsilon_{\mathcal{Q}}) \log(t + \epsilon_{\mathcal{Q}});$$

$$\frac{e^{-\delta_t}}{(1 - \delta_t)^{1-\delta_t}} \leq \exp(-c_t).$$

Thus for $t \in (0, 1)$,

$$\mathbb{P}\left(\|\boldsymbol{M}_{\mathcal{S}} - \boldsymbol{I}\| \leq t \,\middle|\, \boldsymbol{\xi}^{(k)} \in \mathcal{Q} \,\forall k \right) \leq 2s \exp\left(-\frac{c_t(1 - \epsilon_{\mathcal{Q}})N}{s\beta_{\mathcal{Q}}} \right); \tag{61}$$

$$\leq 2s \exp\left(-C_{\mathcal{Q}} \frac{tN}{s\beta_{\mathcal{Q}}} \right), \tag{62}$$

where $C_{\mathcal{Q}}$ is a reasonably small positive constant for most truncations. Via a union bound over the $\binom{P}{s}$ possibilities of subsets of P with cardinality s, it follows that

$$\mathbb{P}\left(\sup_{|\mathcal{S}| \leq s} \lambda_{\max}(\boldsymbol{M}_{\mathcal{S}} - \boldsymbol{I}) \geq t \right) \leq 2s \binom{P}{s} \exp\left(-C_{\mathcal{Q}} \frac{Nt}{s\beta_{\mathcal{Q}}} \right). \tag{63}$$

Remark 5.1. Here we point out the difference between the result from the presented proof and the approach leading to the result in Remark 3.3. The proof here proceeds via a union bound over all possible selections of s columns from the full set of P, which provides servicable bounds for many choices of s and N. However, this leads to an unnecessarily loose bound when considering asymptotically large s and P. Specifically, the tail bound of (61) for an arbitrary subset of s columns is sufficient to bound a number of points that can be used to bound $\lambda_{\max}(\boldsymbol{M}_{\mathcal{S}} - \boldsymbol{I})$ uniformly over choices $\mathcal{S}$ satisfying $|\mathcal{S}| \leq s$. These arguments [39–41] yield a version of (63) which omits the $\binom{P}{s}$ scaling, such that the probability scales linearly in s up to logarithmic factors.

The specifics of these arguments are not repeated here. The interested reader is directed to Section 8 of [39] for analysis of bounded orthonormal systems that translate to the setup presented here. To account for the generalization of independence structure from independent rows to independent sets of rows, (61) demonstrates the appropriateness of the definition in (26) to generalize the coherence of the orthonormal system from that given by (17).

Recalling that

$$\sup_{|\mathcal{S}| \leq s} \lambda_{\max}(\boldsymbol{M}_{\mathcal{S}} - \boldsymbol{I}) = \delta_s,$$

gives

$$\mathbb{P}(\delta_s \geq t) \leq \exp\left(-C_{\mathcal{Q}} \frac{Nt}{s\beta_{\mathcal{Q}}} + \log\left(2s\binom{P}{s} \right) \right). \tag{64}$$

We assume that having any sample $\boldsymbol{\xi}^{(k)} \in \mathcal{Q}^c$ leads to an arbitrarily large $\lambda_{\max}$, hence yielding the bound,

$$\mathbb{P}(\delta_s \geq t) \leq 1 - \mathbb{P}(\mathcal{Q})^N + \exp\left(-C_{\mathcal{Q}} \frac{Nt}{s\beta_{\mathcal{Q}}} + \log\left(2s\binom{P}{s} \right) \right). \tag{65}$$

Using the relation, $\mathbb{P}(\delta_s < t) = 1 - \mathbb{P}(\delta_s \geq t)$, gives that

$$\mathbb{P}(\delta_s < t) \geq \mathbb{P}(\mathcal{Q})^N - \exp\left(-C_{\mathcal{Q}} \frac{Nt}{s\beta_{\mathcal{Q}}} + \log\left(2s\binom{P}{s} \right) \right). \tag{66}$$

Using the relation that

$$\binom{P}{s} \leq \left(\frac{eP}{s} \right)^s,$$

it follows that

$$2s\binom{P}{s} \leq 2se^s\left(\frac{P}{s} \right)^s; \tag{67}$$

$$\log\left(2s\binom{P}{s} \right) \leq \log(2s) + s + s\log(P/s), \tag{68}$$

which completes the proof. Remark 3.2 follows from (66), taking $s = P$ and using that $\binom{P}{P} = 1$. □

We note that Corollary 1 follows from Theorem 3.3, by substituting $p_\star$ for $\mathbb{P}(\delta_s < t)$; substituting $\delta_\star$ for t; and performing some algebraic manipulation.

5.4. Theorem 3.5

The proof of Theorem 3.5 relies on the properties of the measurement matrices $\tilde{\boldsymbol{\Psi}}$ and $\boldsymbol{\Psi}$ themselves. In an intuitive sense, the results follow from the two matrices having similar properties, but the gradient-enhanced matrix having more rows, yielding better conditioned Gramians.

Proof. We begin by showing **R1** for the one-dimensional case. Note that for arbitrary ξ, and i,

$$\frac{|\psi_i(\xi)|^2 + i|\psi_{i-1}(\xi)|^2}{1+i} \leq \max\left\{|\psi_i(\xi)|^2, |\psi_{i-1}(\xi)|^2\right\}$$

and that equality can only hold if $\psi_i(\xi) = \psi_{i-1}(\xi)$ for some i, which is an event that occurs with probability zero. As this inequality holds for all i and ξ, being strict for almost all ξ, **R1** follows for the one-dimensional case. The d-dimensional analogue follows as the d-dimensional polynomials are tensor-products of the one-dimensional polynomials.

To show **R2**, note that up to an invertible post-multiplication, $\boldsymbol{\Psi}$ is a submatrix of $\tilde{\boldsymbol{\Psi}}$, and thus $\mathcal{N}(\tilde{\boldsymbol{\Psi}}\boldsymbol{W}^{-1}) \subset \mathcal{N}(\boldsymbol{\Psi})$. Additionally, we notice that $\boldsymbol{\Psi}$, and $\tilde{\boldsymbol{\Psi}}$ are almost surely full rank matrices.

We next show **R3** in the case of 1-dimensional Hermite polynomials. Here subscripts of matrices refer to the column corresponding to that polynomial order. We have by the normalization (24) and (45) that

$$(\tilde{\boldsymbol{\Psi}}_i, \tilde{\boldsymbol{\Psi}}_j) = \frac{(\boldsymbol{\Psi}_i, \boldsymbol{\Psi}_j) + \sqrt{ij}(\boldsymbol{\Psi}_{i-1}, \boldsymbol{\Psi}_{j-1})}{\sqrt{(1+i)(1+j)}}. \tag{69}$$

It follows that,

$$|(\tilde{\boldsymbol{\Psi}}_i, \tilde{\boldsymbol{\Psi}}_j)| \leq \frac{|(\boldsymbol{\Psi}_i, \boldsymbol{\Psi}_j)| + \sqrt{ij}|(\boldsymbol{\Psi}_{i-1}, \boldsymbol{\Psi}_{j-1})|}{\sqrt{(1+i)(1+j)}}.$$

Applying a supremum,

$$\sup_{i \neq j}|(\tilde{\boldsymbol{\Psi}}_i, \tilde{\boldsymbol{\Psi}}_j)| \leq \sup_{i \neq j} \frac{|(\boldsymbol{\Psi}_i, \boldsymbol{\Psi}_j)| + \sqrt{ij}|(\boldsymbol{\Psi}_{i-1}, \boldsymbol{\Psi}_{j-1})|}{\sqrt{(1+i)(1+j)}},$$

$$\leq \sup_{i \neq j}|(\boldsymbol{\Psi}_i, \boldsymbol{\Psi}_j)|\frac{1 + \sqrt{ij}}{\sqrt{(1+i)(1+j)}},$$

where we have used the inequality

$$\frac{1 + \sqrt{(i-1)(j-1)}}{\sqrt{(1+(i-1))(1+(j-1))}} \leq \frac{1 + \sqrt{ij}}{\sqrt{(1+i)(1+j)}}. \tag{70}$$

This shows the result for the one-dimensional case. The d-dimensional case leads to a decomposition in (69) with inner products of lower order along each dimension, and the inequality in (70) is replaced by d similar inequalities. $\square$

6. Conclusion

Within the context of compressive sampling of sparse polynomial chaos expansions, we investigated ℓ_1-minimization when derivative information of a quantity of interest (QoI) is present. We provided analyses on gradient-enhanced ℓ_1-minimization for Hermite polynomial chaos, in which we showed that, for a given normalization, including derivative information will not reduce the stability of the ℓ_1-minimization problem. Further, we identified a coherence parameter that we used to bound the associated Restricted Isometry Constant, a useful and well-studied measure of the stability for solutions recovered by ℓ_1-minimization as in the context used here.

Furthermore, we observed improved solution accuracy from gradient-enhanced ℓ_1-minimization in three numerical examples: Manufactured polynomials; an elliptic equation; and a plane Poiseuille flow with random boundaries. Consistently, gradient-enhanced ℓ_1-minimization was seen to improve the quality of solution recovery at the same computational cost, or equivalently achieve the same solution quality at a reduced computational cost. As the QoI derivatives are often more sensitive to discretization errors than the QoI itself, so too is the accuracy of the solution obtained by the gradient-enhanced ℓ_1-minimization. This was empirically observed in the second numerical example considered, thereby suggesting high accuracy requirements on derivative calculations for the gradient-enhanced ℓ_1-minimization to be most effective.

Acknowledgements

References

[1] R. Ghanem, P. Spanos, Stochastic Finite Elements: A Spectral Approach, Springer-Verlag, 1991.
[2] O.P. Le 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] D. Xiu, G.M. Karniadakis, The Wiener–Askey polynomial chaos for stochastic differential equations, SIAM J. Sci. Comput. 24 (2) (2002) 619–644.
[5] E.J. Candès, J. Romberg, T. Tao, Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information, IEEE Trans. Inf. Theory 52 (2) (2006) 489–509.
[6] D.L. Donoho, Compressed sensing, IEEE Trans. Inf. Theory 52 (4) (2006) 1289–1306.
[7] A. Doostan, H. Owhadi, A non-adapted sparse approximation of PDEs with stochastic inputs, J. Comput. Phys. 230 (2011) 3015–3034.
[8] Ji Peng, Jerrad Hampton, Alireza Doostan, A weighted ℓ_1-minimization approach for sparse polynomial chaos expansions, J. Comput. Phys. 267 (2014) 92–111.
[9] J. Hampton, A. Doostan, Compressive sampling of polynomial chaos expansions: convergence analysis and sampling strategies, J. Comput. Phys. 280 (2015) 363–386.
[10] M.T. Reagan, H.N. Najm, R.G. Ghanem, O.M. Knio, Uncertainty quantification in reacting-flow simulations through non-intrusive spectral projection, Combust. Flame 132 (3) (2003) 545–555.
[11] L. Mathelin, 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.
[12] D. Xiu, J.S. Hesthaven, High-order collocation methods for differential equations with random inputs, SIAM J. Sci. Comput. 27 (3) (2005) 1118–1139.
[13] P.G. Constantine, M. Eldred, E. Phipps, Sparse pseudospectral approximation method, Comput. Methods Appl. Mech. Eng. 229 (2012) 1–12.
[14] S. Hosder, R.W. 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.
[15] G. Migliorati, Fabio Nobile, Erik von Schwerin, Raúl Tempone, Approximation of quantities of interest in stochastic PDEs by the random discrete ℓ^2 projection on polynomial spaces, SIAM J. Sci. Comput. 35 (3) (2013) A1440–A1460.
[16] J. Hampton, A. Doostan, Coherence motivated sampling and convergence analysis of least squares polynomial chaos regression, Comput. Methods Appl. Mech. Eng. (2015) 73–97.
[17] A. Doostan, H. Owhadi, A. Lashgari, G. Iaccarino, Non-adapted sparse approximation of PDEs with stochastic inputs, Technical report, Annual research brief, Center for Turbulence Research, Stanford University, 2009.
[18] L. Yan, L. Guo, D. Xiu, Stochastic collocation algorithms using ℓ_1-minimization, Int. J. Uncertain. Quantificat. 2 (3) (2012).
[19] X. Yang, G.E. Karniadakis, Reweighted ℓ_1 minimization method for stochastic elliptic differential equations, J. Comput. Phys. 248 (2013) 87–108.
[20] B. Jones, N. Parrish, A. Doostan, Postmaneuver collision probability estimation using sparse polynomial chaos expansions, J. Guid. Control Dyn. (2015) 1–13.
[21] J.D. Jakeman, M.S. Eldred, K. Sargsyan, Enhancing ℓ_1-minimization estimates of polynomial chaos expansions using basis selection, J. Comput. Phys. 289 (2015) 18–34.
[22] Sylvain Arlot, Alain Celisse, A survey of cross-validation procedures for model selection, Stat. Surv. 4 (2010) 40–79.
[23] O. Roderick, M. Anitescu, P. Fischer, Polynomial regression approaches using derivative information for uncertainty quantification, Nucl. Sci. Eng. 162 (2) (2010) 122–139.
[24] A.K. Alekseev, I.M. Navon, M.E. Zelentsov, The estimation of functional uncertainty using polynomial chaos and adjoint equations, Int. J. Numer. Methods Fluids 67 (3) (2011) 328–341.
[25] Y. Li, M. Anitescu, O. Roderick, F. Hickernell, Orthogonal bases for polynomial regression with derivative information in uncertainty quantification, Int. J. Uncertain. Quantificat. 1 (4) (2011).
[26] J.H. de Baar, Thomas P. Scholcz, Clemens V. Verhoosel, Richard P. Dwight, Alexander H. van Zuijlen, Hester Bijl, Efficient uncertainty quantification with gradient-enhanced kriging: applications in FSI, in: Proc. of the European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS 2012), Vienna, Austria, 2012, pp. 10–14.
[27] Brian Lockwood, Dimitri Mavriplis, Gradient-based methods for uncertainty quantification in hypersonic flows, Comput. Fluids: International Workshop on Future of {CFD} and Aerospace Sciences 85 (2013) 27–38.
[28] Holger Rauhut, Rachel Ward, Interpolation via weighted minimization, Appl. Comput. Harmon. Anal. (2015).
[29] Ben Adcock, Infinite-dimensional ℓ_1 minimization and function approximation from pointwise data, arXiv:1503.02352, 2015.
[30] Vadim Komkov, Kyung K. Choi, Edward J. Haug, Design Sensitivity Analysis of Structural Systems, vol. 177, Academic Press, 1986.
[31] L.B. Rall, Automatic Differentiation: Techniques and Applications, vol. 120, Springer, Berlin, 1981.
[32] Gary Tang, Methods for high dimensional uncertainty quantification: regularization, sensitivity analysis, and derivative enhancement, PhD thesis, Stanford University, 2013.
[33] E.J. Candès, The restricted isometry property and its implications for compressed sensing, C. R. Math. 346 (9) (2008) 589–592.
[34] R.A. Askey, W.J. Arthur, Some Basic Hypergeometric Orthogonal Polynomials that Generalize Jacobi Polynomials, vol. 319, AMS, Providence, RI, 1985.
[35] G. Szegö, Orthogonal Polynomials, American Mathematical Society, 1939.
[36] Holger Rauhut, Rachel Ward, Sparse Legendre expansions via l1-minimization, J. Approx. Theory 164 (5) (May 2012) 517–533.
[37] E.J. Candès, Y. Plan, A probabilistic and ripless theory of compressed sensing, IEEE Trans. Inf. Theory 57 (11) (2010) 7235–7254.
[38] Albert Cohen, Mark A. Davenport, Dany Leviatan, On the stability and accuracy of least squares approximations, Found. Comput. Math. 13 (5) (2013) 819–834.
[39] Holger Rauhut, Compressive sensing and structured random matrices, 2009.
[40] Felix Krahmer, Shahar Mendelson, Holger Rauhut, Suprema of chaos processes and the restricted isometry property, Commun. Pure Appl. Math. 67 (11) (2014) 1877–1904.
[41] Mark Rudelson, Roman Vershynin, On sparse reconstruction from Fourier and Gaussian measurements, Commun. Pure Appl. Math. 61 (8) (2008) 1025–1045.
[42] A. Juditsky, A. Nemirovski, Accuracy guarantees for ℓ_1-recovery, IEEE Trans. Inf. Theory 57 (2011) 7818–7839.
[43] A. Juditsky, A. Nemirovski, On verifiable sufficient conditions for sparse signal recovery via ℓ_1 minimization, Math. Program. 127 (1) (2011) 57–88.
[44] A. Logg, K. Mardal, G. Wells, Automated Solution of Differential Equations by the Finite Element Method, Springer, Berlin, Heidelberg, 2012.
[45] J.F. Sykes, J.L. Wilson, R.W. Andrews, Sensitivity analysis for steady state groundwater flow using adjoint operators, Water Resour. Res. 21 (3) (1985) 359–371.
[46] Yang Cao, Shengtai Li, Linda Petzold, Radu Serban, Adjoint sensitivity analysis for differential–algebraic equations: the adjoint DAE system and its numerical solution, SIAM J. Sci. Comput. 24 (3) (2003) 1076–1089.
[47] Michael B. Giles, Niles A. Pierce, An introduction to the adjoint approach to design, Flow Turbul. Combust. 65 (3–4) (2000) 393–415.
[48] Arthur Earl Bryson, Applied Optimal Control: Optimization, Estimation and Control, CRC Press, 1975.
[49] Emmanuel Laporte, Patrick Le Tallec, Numerical Methods in Sensitivity Analysis and Shape Optimization, Springer Science & Business Media, 2003.

[50] Antony Jameson, Aerodynamic shape optimization using the adjoint method, Lectures at the Von Karman Institute, Brussels, 2003.
[51] J.A. Tropp, User-friendly tail bounds for sums of random matrices, Found. Comput. Math. 12 (4) (2012) 389–434.
[52] Richard Baraniuk, Mark Davenport, Ronald DeVore, Michael Wakin, A simple proof of the restricted isometry property for random matrices, Constr. Approx. 28 (3) (2008) 253–263.