

Reconstructing high-dimensional Hilbert-valued functions via compressed sensing

Nick Dexter, Hoang Tran, and Clayton Webster

Abstract—We present and analyze a novel sparse polynomial technique for approximating high-dimensional Hilbert-valued functions, with application to parameterized partial differential equations (PDEs) with deterministic and stochastic inputs. Our theoretical framework treats the function approximation problem as a joint sparse recovery problem, where the set of jointly sparse vectors is possibly infinite. To achieve the simultaneous reconstruction of Hilbert-valued functions in both parametric domain and Hilbert space, we propose a novel mixed-norm based ℓ_1 regularization method that exploits both energy and sparsity. Our approach requires extensions of concepts such as the restricted isometry and null space properties, allowing us to prove recovery guarantees for sparse Hilbert-valued function reconstruction. We complement the enclosed theory with an algorithm for Hilbert-valued recovery, based on standard forward-backward algorithm, meanwhile establishing its strong convergence in the considered infinite-dimensional setting. Finally, we demonstrate the minimal sample complexity requirements of our approach, relative to other popular methods, with numerical experiments approximating the solutions of high-dimensional parameterized elliptic PDEs.

Index Terms—High-dimensional approximation, compressed sensing, Hilbert-valued functions, parametric PDEs, bounded orthonormal systems, forward-backward iterations.

I. INTRODUCTION

Underlying many successful applications of compressed sensing to problems in applied mathematics and the physical sciences is the fact that, for many practical problems, the target to be reconstructed possesses sufficient sparsity in order to enable unique solutions from systems that would otherwise be ill-posed. In the basic compressed sensing problem, the signal is an unknown vector $\mathbf{c} \in \mathbb{R}^N$, and the sensing process yields a measurement vector $\mathbf{u} \in \mathbb{R}^m$ that is formed by the product of \mathbf{c} with a sensing matrix, i.e., $\mathbf{u} = \mathbf{A}\mathbf{c}$, where $\mathbf{A} \in \mathbb{R}^{m \times N}$. The key observation is that when the signal \mathbf{c} is sufficiently sparse, it can be uniquely determined from an underdetermined set of measurements ($m \ll N$), provided \mathbf{A}

This material is based upon work supported in part by: the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under contracts and awards ERKJ314, ERKJ331, ERKJ345, and Scientific Discovery through Advanced Computing (SciDAC) program through the FASTMath Institute under Contract No. DE-AC02-05CH11231; and by the Laboratory Directed Research and Development program at the Oak Ridge National Laboratory, which is operated by UT-Battelle, LLC., for the U.S. Department of Energy under contract DE-AC05-00OR22725.

Department of Mathematics, Simon Fraser University. Email: nicholas_dexter@sfu.ca

Computer Science and Mathematics Division, Oak Ridge National Laboratory. Email: tranha@ornl.gov

Department of Mathematics, University of Tennessee, Knoxville, and Computer Science and Mathematics Division, Oak Ridge National Laboratory. Email: webstercg@math.utk.edu

satisfies certain additional properties. To overcome the NP-hardness of directly finding the sparsest \mathbf{c} consistent with a given set of measurements, various greedy and convex relaxation strategies have been proposed and demonstrated, both empirically and theoretically, to have good reconstruction performance in a range of settings.

In the parameterized PDE literature, one typically seeks to approximate the parameter-to-solution map $\mathbf{y} \mapsto u(\mathbf{y})$, taking values in a Hilbert space \mathcal{V} , by a truncation of its orthonormal expansion, i.e.,

$$u = \sum_{\nu \in \mathcal{F}} \mathbf{c}_\nu \Psi_\nu \approx \sum_{\nu \in \mathcal{J}} \mathbf{c}_\nu \Psi_\nu =: u_{\mathcal{J}}. \quad (1)$$

Here, $\mathbf{y} \in \mathcal{U}$ a high-dimensional tensor product domain, \mathcal{J} is a finite multi-index set of cardinality $N = \#(\mathcal{J})$ with $\mathcal{J} \subset \mathcal{F} := \mathbb{N}_0^d$, $(\Psi_\nu)_{\nu \in \mathcal{J}}$ is an $L^2(\mathcal{U}, d\rho)$ -orthonormal basis, and $\mathbf{c} = (\mathbf{c}_\nu)_{\nu \in \mathcal{J}} \in \mathcal{V}^N$ are the Hilbert-valued coefficients to be computed. Generally, one selects \mathcal{J} to be a set large enough to ensure that

$$\|u - u_{\mathcal{J}}\|_{L^2}^2 = \|u_{\mathcal{J}^c}\|_{L^2}^2 = \sum_{\nu \in \mathcal{F} \setminus \mathcal{J}} \|\mathbf{c}_\nu\|_{\mathcal{V}}^2 \quad (2)$$

is minimal. However, this can lead to a less efficient approximation if \mathcal{J} is not chosen carefully.

Over the course of the last decade, a series of works in the parameterized PDEs community (see [4] and the reference therein) have demonstrated that, under reasonable assumptions on the input data to the PDEs, the solutions are compressible, and hence well-represented by sparse expansions in given orthonormal systems. In other words, the solution vector \mathbf{c} from (1) is sparse, and accurate reconstructions of the most important components of \mathbf{c} are enough for satisfactory approximations of the PDE solutions. The ability of compressed sensing (CS) to exploit sparsity and allow far fewer samples than traditional approaches (e.g., Monte-Carlo, projection, interpolation) makes it a promising tool for such reconstruction problems. Therefore it is no surprise that CS-based polynomial approximation has attracted growing interest in the area of high-dimensional parameterized PDEs in recent years, [8], [12], [13], [15]–[17], [21].

However, there has been a critical mismatch between standard CS techniques and the problem of reconstructing Hilbert-valued solutions to parameterized PDEs: these methods do not enable direct recovery of the vector $(\mathbf{c}_\nu)_{\nu \in \mathcal{J}} \in \mathcal{V}^N$, i.e., a vector with Hilbert-valued components. Instead, CS-based polynomial approximation methods only allow the recovery of real or complex sparse vectors. Hence, in the context of

parameterized PDEs, standard CS-based polynomial approximation methods do not yield approximations of the entire solution map $\mathbf{y} \mapsto u(\cdot, \mathbf{y}) \in \mathcal{V}$, but only functionals of the solution, i.e., maps of the form $\mathbf{y} \mapsto G(u(\mathbf{y})) \in \mathbb{R}$, where $G : \mathcal{V} \rightarrow \mathbb{R}$ is a functional of u . In particular, many of the existing works (cited above) perform *point-wise reconstruction* of the solution, i.e., reconstruction of the map $\mathbf{y} \mapsto G_{x^*}(u(\mathbf{y}))$ with $G_{x^*}(u(\mathbf{y})) = u(x^*, \mathbf{y})$ for a fixed $x^* \in D$, with D the physical domain of the PDE.

In this paper, we present an overview of latest progress in addressing this mismatch by developing a novel sparse recovery technique which enables direct reconstruction of Hilbert-valued vectors, [5], [6]. The theoretical framework, based on the problem of joint-sparse recovery, treats the recovery problem as a matrix recovery problem in which each row of the solution matrix may have infinitely many terms corresponding to the coefficients of a Hilbert-valued function in a given basis. Our regularization is performed with respect to a mixed norm $\|\cdot\|_{\mathcal{V},1}$, which is defined to be the ℓ_1 norm of the vector $(\|\mathbf{c}_\nu\|_{\mathcal{V}})_{\nu \in \mathcal{J}}$. Often, the decay of the polynomial coefficients and tail bounds are estimated in global energy norms rather than pointwise over the physical domain D [4]. As we will show, our choice of regularization enables us to prove convergence rates for global approximations in the same energy norms, with explicit computable coefficients. We present the theoretical support of this strategy via several extensions of compressed sensing concepts such as the restricted isometry property (RIP) and null space property (NSP) to the Hilbert-valued setting.

Our regularization problem is solved with a forward-backward splitting approach. We present a new strong convergence result for forward-backward splitting in a joint sparse recovery scenario where we assume neither the strict convexity of the fidelity functions nor the finite cardinality of the set of signals to construct. We note that most of similar strong convergence results along the line need at least one of these conditions. In relation to parameterized PDEs problem, this result means that forward-backward splitting strongly converges even before discretization in (infinite-dimensional) Hilbert space is introduced.

The rest of the paper proceeds as follows. In Section II, we introduce our approach of sparse regularization and provide a brief comparison to the problem of joint-sparse recovery. In Section III, we present theoretical guarantees for the direct reconstruction of Hilbert-valued functions through mixed norm regularization, as well as convergence estimates in case of parameterized PDEs. Section IV introduces a version of forward-backward algorithm adapted for Hilbert-valued recovery and provides its strong convergence result. Section V presents numerical results on applying the sparse regularization method to the solution of a parameterized elliptic PDE. Finally, Section VI concludes with some remarks on our approach and applications to a wider array of problems.

II. SPARSE REGULARIZATION FOR PARAMETERIZED PDES

The sparse polynomial techniques proposed in this work are applicable to general parameterized PDE problems of the

form: find $u(\cdot, \mathbf{y}) : \bar{D} \times \mathcal{U} \rightarrow \mathbb{R}$ for all $\mathbf{y} \in \mathcal{U}$ such that

$$\mathcal{L}(u(\cdot, \mathbf{y}), \mathbf{y}) = 0, \quad \text{in } D, \quad (3)$$

where \mathcal{L} is a differential operator defined on a spatio-temporal domain D . Our development to the general CS/joint-sparse polynomial approximation problem can be formulated as follows. One first generates m samples $\mathbf{y}_1, \dots, \mathbf{y}_m$ in \mathcal{U} independently from the orthogonalization measure ϱ associated with $(\Psi_\nu)_{\nu \in \mathcal{J}}$, for instance, uniform samples for the Legendre basis and Chebyshev samples for Chebyshev basis, see [9, Chapter 12], and solves the equation (3) at these samples to form the normalized output vector $\mathbf{u} := (u(\mathbf{y}_1), \dots, u(\mathbf{y}_m))/\sqrt{m}$ as well as the normalized sampling matrix $\mathbf{A} := (\Psi_\nu(\mathbf{y}_i)/\sqrt{m})_{1 \leq i \leq m, \nu \in \mathcal{J}}$. Taking account that the true unknown coefficient $\mathbf{c} = (\mathbf{c}_\nu)_{\nu \in \mathcal{J}}$ approximately solves the linear system

$$\mathbf{u} = \mathbf{A}\mathbf{z}, \quad \mathbf{z} \in \mathcal{V}^N, \quad (4)$$

and further, \mathbf{c} is compressible [4], it is reasonable to approximate \mathbf{c} by $\mathbf{c}^\#$, the solution to the following problem

$$\min_{\mathbf{z} \in \mathcal{V}^N} \|\mathbf{z}\|_{\mathcal{V},1} \quad \text{subject to} \quad \|\mathbf{A}\mathbf{z} - \mathbf{u}\|_{\mathcal{V},2} \leq \frac{\eta}{\sqrt{m}}, \quad (5)$$

where η relates to an estimate of the tail (2); or the equivalent unconstrained convex minimization:

$$\min_{\mathbf{z} \in \mathcal{V}^N} \|\mathbf{z}\|_{\mathcal{V},1} + \frac{\mu}{2} \|\mathbf{A}\mathbf{z} - \mathbf{u}\|_{\mathcal{V},2}^2, \quad (6)$$

for appropriately chosen $\mu > 0$. Here, the norm $\|\cdot\|_{\mathcal{V},q}$ is defined for $\mathbf{c} \in \mathcal{V}^N$ as $\|\mathbf{c}\|_{\mathcal{V},q} := (\sum_{\nu \in \mathcal{J}} \|\mathbf{c}_\nu\|_{\mathcal{V}}^q)^{1/q}$. This is arguably the most natural extension of the ℓ_1 minimization approach, traditionally for real and complex signal recovery, to the reconstruction of sparse generalized Hilbert-valued vectors. We denote our approach *simultaneous compressed sensing* (SCS). In depth description, analysis and application of SCS for solving parameterized PDEs are provided in [6].

Problem (5) can be related to the joint-sparse basis pursuit denoising problem as follows. Let $(\phi_r)_{r \in \mathbb{N}}$ be an orthonormal basis of \mathcal{V} , then $\mathbf{c}_\nu \in \mathcal{V}$ has unique representation

$$\mathbf{c}_\nu = \sum_{r \in \mathbb{N}} c_{\nu,r} \phi_r, \quad \text{with } c_{\nu,r} \in \mathbb{R}.$$

Each coefficient \mathbf{c}_ν corresponds to an $\mathbb{R}^{1 \times \mathbb{N}}$ vector $(c_{\nu,1}, c_{\nu,2}, \dots, c_{\nu,r}, \dots)^T$, thus, $\mathbf{c} = (\mathbf{c}_\nu)_{\nu \in \mathcal{J}}$ is completely determined by the $\mathbb{R}^{N \times \mathbb{N}}$ matrix $\hat{\mathbf{c}} = (c_{\nu,r})_{\nu \in \mathcal{J}, r \in \mathbb{N}}$. Furthermore,

$$\begin{aligned} \|\mathbf{c}\|_{\mathcal{V},q} &= \left(\sum_{\nu \in \mathcal{J}} \|\mathbf{c}_\nu\|_{\mathcal{V}}^q \right)^{1/q} = \left(\sum_{\nu \in \mathcal{J}} \left(\sum_{r \in \mathbb{N}} |c_{\nu,r}|^2 \right)^{q/2} \right)^{1/q} \\ &\equiv \left(\sum_{\nu \in \mathcal{J}} \|\hat{\mathbf{c}}_\nu\|_2^q \right)^{1/q} = \|\hat{\mathbf{c}}\|_{2,q}, \end{aligned}$$

where the matrix norm $\|\cdot\|_{p,q}$ is defined as the ℓ_q norm of the vector $(\|\hat{\mathbf{c}}_\nu\|_p)_{\nu \in \mathcal{J}}$, implying the equivalence of (5) with the infinite-dimensional joint-sparse recovery problem. In practice, one needs to employ a discretization over \mathcal{V} to be able to numerically solve (5) or (6). Any preferred method of spatial/temporal discretization may be used, e.g., the finite

element, difference, and volume methods. However, our strong convergence result of forward-backward splitting in Section IV is applicable to (6) without any discretization.

III. ERROR ESTIMATES FOR HILBERT-VALUED RECOVERY

Straightforward extensions of concepts and results from compressed sensing and joint-sparse recovery can be made to ensure uniform recovery of Hilbert-valued signals via $\ell_{\mathcal{V},1}$ -relaxation. Well-known concepts such as the null space property (NSP) and restricted isometry property (RIP) have Hilbert-valued counterparts. In this section, we review Hilbert-valued versions of the NSP and RIP to guarantee uniform recovery for the mixed norm regularization in the presence of noise or sparsity defects. We note that the extension of these concepts to the Hilbert-valued recovery setting does not require finite-dimensionality of the Hilbert space \mathcal{V} . For more detailed discussion and complete proofs of the results in this section, we refer the interested readers to [6].

First, we define a Hilbert-valued version of the ℓ_2 -robust NSP, which guarantees the reconstruction of vectors $\mathbf{c} \in \mathcal{V}^N$ (up to s largest components and up to a noise level).

Definition 1 ($\ell_{\mathcal{V},2}$ -robust null space property). *The matrix $\mathbf{A} \in \mathbb{R}^{m \times N}$ is said to satisfy the $\ell_{\mathcal{V},2}$ -robust null space property of order s with constants $0 < \rho < 1$ and $\tau > 0$ if*

$$\|\mathbf{z}_S\|_{\mathcal{V},2} \leq \frac{\rho}{\sqrt{s}} \|\mathbf{z}_{S^c}\|_{\mathcal{V},1} + \tau \|\mathbf{A}\mathbf{z}\|_{\mathcal{V},2} \quad (7)$$

$$\forall \mathbf{z} \in \mathcal{V}^N, \forall S \subset [N] \text{ with } \#(S) \leq s.$$

An RIP-type condition is required to quantify the sample complexity of solving (5) to a prescribed accuracy. The following result establishes the implication of the $\ell_{\mathcal{V},2}$ -robust NSP from the standard RIP.

Proposition 1. *Suppose that \mathcal{V} is a separable Hilbert space, and that the matrix $\mathbf{A} \in \mathbb{R}^{m \times N}$ satisfies the RIP, that is*

$$(1 - \delta_{2s}) \|\mathbf{z}\|_2^2 \leq \|\mathbf{A}\mathbf{z}\|_2^2 \leq (1 + \delta_{2s}) \|\mathbf{z}\|_2^2, \quad (8)$$

$$\forall \mathbf{z} \in \mathbb{R}^N, \mathbf{z} \text{ } 2s\text{-sparse},$$

with $\delta_{2s} < \frac{4}{\sqrt{41}}$. Then, \mathbf{A} satisfies the $\ell_{\mathcal{V},2}$ -robust NSP of order s with constants $0 < \rho < 1$ and $\tau > 0$ depending only on δ_{2s} .

Proposition 1 implies that sample complexity results for solving the standard ℓ_1 problem hold for the mixed norm problem (5) as well.

In the parameterized PDE context, with error estimates as in [19, Theorem 2] for quasi-optimal approximations, we are able to provide convergence rates for approximations to parameterized PDEs obtained through the mixed norm regularization. First, we assume that the solution u has parametric expansion with coefficients $(\mathbf{c}_\nu)_{\nu \in \mathcal{F}}$ as in (1) satisfying $\|\mathbf{c}_\nu\|_{\mathcal{V}} \lesssim e^{-b(\nu)}$ for every $\nu \in \mathcal{F}$, with $b(\nu)$ obeying [19, Assumption 3]. For brevity, we do not detail that assumption herein, but remark that it is satisfied by most parameterized PDE models we are aware of.

Theorem 1. *For any $\varepsilon > 0$, assume that the solution u to (3) with parametric expansion (1) has coefficients $(\mathbf{c}_\nu)_{\nu \in \mathcal{F}}$*

satisfying $\|\mathbf{c}_\nu\|_{\mathcal{V}} \leq e^{-b(\nu)}$ for every $\nu \in \mathcal{F}$ with $b : [0, \infty)^d \rightarrow \mathbb{R}$ also satisfying [19, Assumption 3]. Denote by \mathcal{J}_s the set of indices corresponding to the s largest bounds of the sequence $(e^{-b(\nu)})_{\nu \in \mathcal{F}}$. Let \mathcal{J} be such that $\mathcal{J}_s \subseteq \mathcal{J}$, and assume that the number of samples $m \in \mathbb{N}$ satisfies

$$m \geq C\Theta^2 s \max\{\log^2(\Theta^2 s) \log(N), \log(\Theta^2 s) \log(\log(\Theta^2 s) N^{\log(s)})\}, \quad (9)$$

with $\Theta = \sup_{\nu \in \mathcal{J}} \|\Psi_\nu\|_{L^\infty(\mathcal{U})}$ and $N = \#(\mathcal{J})$. Then with probability $1 - N^{-\log(s)}$, the solution $\mathbf{c}^\#$ of

$$\text{minimize}_{\mathbf{z} \in \mathcal{V}^N} \|\mathbf{z}\|_{\mathcal{V},1} \quad \text{s.t.} \quad \|\mathbf{A}\mathbf{z} - \mathbf{u}\|_{\mathcal{V},2} \leq \frac{\eta}{\sqrt{m}} \quad (10)$$

approximates u with asymptotic error

$$\left\| u - \sum_{\nu \in \mathcal{J}} \mathbf{c}_\nu^\# \Psi_\nu \right\|_{L^2_\varepsilon(\mathcal{U}; \mathcal{V})} \leq C_\varepsilon \sqrt{s} \exp \left[- \left(\frac{\kappa s}{(1 + \varepsilon)} \right)^{1/d} \right], \quad (11)$$

where $\kappa, C_\varepsilon > 0$ are independent of s .

IV. FORWARD-BACKWARD ALGORITHM FOR HILBERT-VALUED RECOVERY

In this section, we present a forward-backward splitting algorithm for solving problem (6) over the real Hilbert space \mathcal{V}^N . Assuming $\mu = 1$ for simplicity, define T_1 and T_2 to be the subdifferential and gradient of the $\|\cdot\|_{\mathcal{V},1}$ and $\frac{1}{2}\|\mathbf{A}(\cdot) - \mathbf{u}\|_{\mathcal{V},2}^2$ parts of (6), respectively, and $T = T_1 + T_2$. The algorithm can be derived as follows. Given $\tau > 0$, we have

$$\mathbf{0} \in T(\mathbf{x}) \iff \mathbf{x} = (I + \tau T_1)^{-1} (I - \tau T_2) \mathbf{x}. \quad (12)$$

The last identity in (12) leads to the *forward-backward splitting* algorithm: given initial guess $\mathbf{x}^0 \in \mathcal{H} := \{\mathbf{x} \in \mathcal{V}^N : \|\mathbf{x}\|_{\mathcal{V},2} < \infty\}$, compute

$$\mathbf{x}^{k+1} = (I + \tau T_1)^{-1} (I - \tau T_2) \mathbf{x}^k, \quad (13)$$

where \mathbf{x}^k denotes the approximation at k -th iterate. Letting $J_\tau := (I + \tau T_1)^{-1}$, $G_\tau := (I - \tau T_2)$, and $S_\tau := J_\tau \circ G_\tau$, then (13) can be written as $\mathbf{x}^{k+1} = S_\tau(\mathbf{x}^k)$. A straightforward derivation yields:

$$G_\tau(\mathbf{x}) = \mathbf{x} - \tau \mathbf{A}^*(\mathbf{A}\mathbf{x} - \mathbf{u}), \quad (14)$$

$$J_{\tau,j}(\mathbf{x}) = \frac{\mathbf{x}_j}{\|\mathbf{x}_j\|_2} \cdot \max\{\|\mathbf{x}_j\|_2 - \tau, 0\}, \quad 1 \leq j \leq N. \quad (15)$$

One can observe that the forward operator G_τ resembles a step of gradient descent algorithm with stepsize τ . The backward operator J_τ , on the other hand, is a soft thresholding step associated with proximal point method.

Under a standard assumption on the stepsize τ , namely $0 < \tau < 2/\|\mathbf{A}^* \mathbf{A}\|_2$, we can obtain some nonexpansive properties for G_τ and J_τ . In particular, G_τ is nonexpansive, i.e.,

$$\|G_\tau(\mathbf{v}) - G_\tau(\mathbf{w})\|_{2,2} \leq \|\mathbf{v} - \mathbf{w}\|_{2,2} \quad \forall \mathbf{v}, \mathbf{w} \in \mathcal{H}, \quad (16)$$

and J_τ is row-wise firmly nonexpansive, i.e., $\forall \mathbf{v}, \mathbf{w} \in \mathcal{H}$, $\forall j \in [N] := \{1, 2, \dots, N\}$

$$\|J_\tau(\mathbf{v}_j) - J_\tau(\mathbf{w}_j)\|_2^2 \leq \|\mathbf{v}_j - \mathbf{w}_j\|_2^2 - \|(I - J_\tau)\mathbf{v}_j - (I - J_\tau)\mathbf{w}_j\|_2^2. \quad (17)$$

see [2, Chapter 4]. To prove the strong convergence of the forward-backward splitting algorithm for Hilbert-valued recovery in infinite-dimensional setting, we consider certain partition of the index set $[N]$ into two subsets L and E , $L \subset (\text{supp}(\mathbf{x}^*))^c$ and $E \supset \text{supp}(\mathbf{x}^*)$ where \mathbf{x}^* is a solution to (6), as they require different treatments, see [5], [11]. With properties (16)-(17) in hand, our result is obtained in three steps:

- 1) Using a partitioning technique from [11], establish finite convergence $\mathbf{x}_j^k \rightarrow \mathbf{x}_j^*$ for every $j \in L$.
- 2) Establish an angular convergence result on the extended support set E .
- 3) Combine the known general weak convergence of the forward-backward algorithm, see, e.g., [2], with our angular convergence result to obtain convergence in norm on E , and hence strong convergence since $L \cup E = [N]$.

For more details on the above approach, see [5]. The established result can then be summarized as follows.

Theorem 2. *Let $0 < \tau < 2/\|A^*A\|_2$ and (\mathbf{x}^k) be the sequence generated by the forward-backward iterations $\mathbf{x}^{k+1} = S_\tau(\mathbf{x}^k)$ starting from any $\mathbf{x}^0 \in \mathcal{H}$. Then (\mathbf{x}^k) converges strongly to an element $\mathbf{x}^* \in \mathcal{V}^N$ solving (6).*

V. NUMERICAL EXPERIMENTS ON PARAMETERIZED ELLIPTIC PDE MODELS

In this section we present numerical experiments demonstrating the efficiency of the proposed approach in approximating the solution of the following parameterized elliptic PDE: for all $\mathbf{y} \in \mathcal{U}$, find $u(\cdot, \mathbf{y}) : D \rightarrow \mathbb{R}$ such that

$$\begin{cases} -\nabla \cdot (a(x, \mathbf{y}) \nabla u(x, \mathbf{y})) = f(x) & \text{in } D, \\ u(x, \mathbf{y}) = 0 & \text{on } \partial D. \end{cases} \quad (18)$$

Here $D = [0, 1]^2$, $f \equiv 1$, the diffusion coefficient is given by

$$a(x, \mathbf{y}) = 10 + y_1 \left(\frac{\sqrt{\pi}L}{2} \right)^{1/2} + \sum_{i=2}^d \zeta_i \vartheta_i(x) y_i \quad (19)$$

$$\zeta_i := (\sqrt{\pi}L)^{1/2} \exp\left(-\frac{\left(\lfloor \frac{i}{2} \rfloor \pi L\right)^2}{8}\right), \text{ for } i > 1,$$

$$\vartheta_i(x) := \begin{cases} \sin\left(\lfloor \frac{i}{2} \rfloor \pi x_1 / L_p\right), & \text{if } i \text{ is even,} \\ \cos\left(\lfloor \frac{i}{2} \rfloor \pi x_1 / L_p\right), & \text{if } i \text{ is odd.} \end{cases}$$

or its log transformed version $\log(a(x, \mathbf{y}) - 0.5)$, and y_i are uniform random variables on $(-\sqrt{3}, \sqrt{3})$. Convergence is compared against “highly-enriched” reference sparse grid stochastic collocation approximations which we denote $u_{h,\text{ex}}(x, \mathbf{y})$, see [18]. All approximations (including the enriched reference solution) are computed on fixed finite element meshes \mathcal{T}_h , and our enriched SC approximation is computed using *Clenshaw-Curtis* (CC) abscissas with high level L_{ex} . We compare performance using the relative errors of the expectation and standard deviation in the $H_0^1(D)$ -norm.

In our plots and discussion, we use the following abbreviations. For the SCS method, we use “SCS-TD” to denote approximations obtained using the SCS method with the *total degree* (TD) subspace. For the *stochastic Galerkin* (SG) method, we use “SG-TD” to denote the SG approximation

with TD subspace, see, e.g., [10], [20]. The *stochastic collocation* (SC) method with CC points and doubling growth rule is denoted “SC-CC,” see [14]. The Monte Carlo approximation is denoted “MC.”

We follow convention from [1], identifying *stochastic degrees of freedom* (SDOF) as the number of random sample points m for MC and SCS, and sparse grid points m_L for SC with level L . For SG, we define SDOF to be N , the cardinality of the index set used in construction. We include the SG method only to compare the $L^2(\mathcal{U}, d\varrho)$ -optimal (w.r.t. SG SDOF) error of the Galerkin projection against the error of the sampling-based approximations.

We employ the orthonormal Legendre series for the parametric discretization of the SG and SCS methods. For the MC and SCS methods, the random samples $(\mathbf{y}_i)_{i=1}^m$ are drawn from the uniform distribution ϱ . We average the random sampling results over 24 trials, fixing the initial seed for the pseudorandom number generator on each trial, and then solving each trial’s problem with the same set of m samples when plotting convergence.

Figures 1 & 2 display the achieved results. For highly anisotropic problems, the compressed-sensing based approach is able to naturally detect the underlying anisotropy in refinement. We note that, for the problem under consideration in Figure 2, we are unable to obtain an SG-TD approximation due to the difficulty of solving the nonlinearly coupled SG systems, see [7] for more details. We also expect that incorporation of anisotropic weighting schemes such as those considered in [1] and weighted ℓ_1 -minimization techniques such as those from [3] to improve results for the problems considered. We leave a further study of such improvements to a future work.

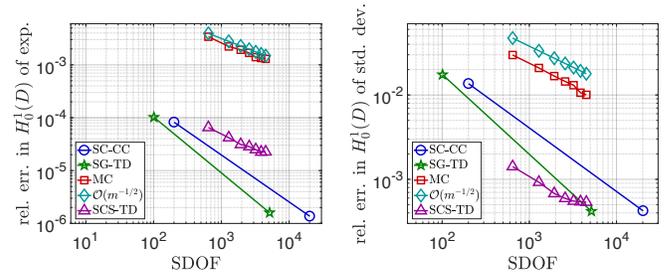


Fig. 1: Comparison of relative errors in expectation (**left**) and standard deviation (**right**) in the $H_0^1(D)$ norm for the $L = 2, 3$ stochastic collocation (SC-CC), $p = 1, 2$ stochastic Galerkin (SG-TD), Monte Carlo (MC), and total degree order $p = 2$ with $N = 5151$ simultaneous compressed sensing (SCS-TD) methods for solving (18) with coefficient (19) and correlation length $L_c = 1/2$ in $d = 100$ dimensions. All approximations computed on a finite element mesh with 206 degrees of freedom corresponding to a maximum mesh size of $h \approx 1/16$.

VI. CONCLUDING REMARKS

In this work we presented an overview of a novel sparse polynomial approximation technique, enabling global recovery of solutions to parameterized PDEs. Our approach, based on

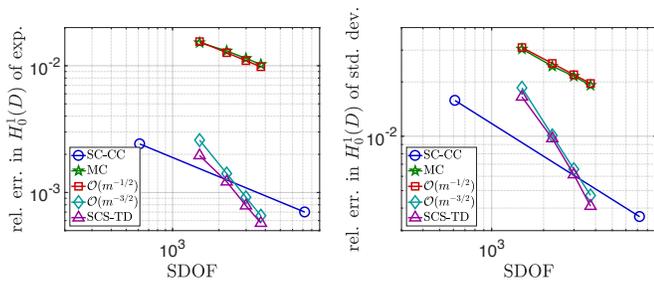


Fig. 2: Comparison of relative errors in expectation (**left**) and standard deviation (**right**) in the $H_0^1(D)$ norm for the $L = 2, 3$ stochastic collocation (SC-CC), Monte Carlo (MC), and total degree order $p = 4$ with $N = 5985$ simultaneous compressed sensing (SCS-TD) methods for solving (18) with coefficient $\log(a(x, \mathbf{y}) - 0.5)$ with $a(x, \mathbf{y})$ from (19) and correlation length $L_c = 1/8$ in $d = 17$ dimensions. All approximations computed on a finite element mesh with 713 degrees of freedom corresponding to a maximum mesh size of $h \approx 1/32$.

extensions of compressed sensing and joint-sparse recovery, treats the solution vector as an element of a tensor product of real Hilbert spaces \mathcal{V} . The key difference in our technique is the use of a mixed norm involving the energy norm of the associated PDE problem and the standard vector ℓ_1 norm. Within this framework, we are able to prove uniform recovery results through straightforward extensions of concepts such as the restricted isometry and null space properties. Moreover, by combining extensions of error estimates for the standard basis pursuit denoising problem with quasi-optimal error estimates for sparse approximation of parameterized PDE systems, we are able to derive sub-exponential convergence of our method. For more details see [6].

We have also presented a modification of the standard forward-backward splitting algorithm for Hilbert-valued recovery. As the considered convex optimization problem is posed over an infinite dimensional Hilbert space, new techniques were needed to establish its convergence properties. By deriving a novel angular convergence result from the firm nonexpansiveness of the soft-thresholding operator, we are able to prove the strong convergence of the algorithm in the considered setting.

Finally, we presented numerical results on the application of our approach to the approximation of solutions to both affinely and non-affinely parameterized elliptic PDEs. We compare our results with those obtained using the stochastic Galerkin, stochastic collocation, and Monte Carlo methods. The achieved results are positive, highlighting a key benefit of compressed sensing-based approaches, namely the ability to detect underlying problem anisotropy in refinement.

REFERENCES

[1] Joakim Bäck, Fabio Nobile, Lorenzo Tamellini, and Raúl Tempone. Stochastic Spectral Galerkin and Collocation Methods for PDEs with Random Coefficients: A Numerical Comparison. In Jan S. Hesthaven and Einar M. Rønquist, editors, *Spectral and High Order Methods for Partial Differential Equations*, volume 76 of *Lecture Notes in*

Computational Science and Engineering, pages 43–62. Springer Berlin Heidelberg, 2011.

[2] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer International Publishing, 2010.

[3] Abdellah Chkifa, Nick Dexter, Hoang Tran, and Clayton G. Webster. Polynomial approximation via compressed sensing of high-dimensional functions on lower sets. *Mathematics of Computation*, 87(311):1415–1450, 2018.

[4] A. Cohen and R. DeVore. Approximation of high-dimensional parametric PDEs. *Acta Numerica*, 24:1–159, 2015.

[5] N. Dexter, H. Tran, and C. G. Webster. On the strong convergence of forward-backward splitting in reconstructing jointly sparse signals. *submitted*, November 2017.

[6] Nick Dexter, Hoang Tran, and Clayton Webster. A mixed ℓ_1 regularization approach for sparse simultaneous approximation of parameterized PDEs. *submitted*, 2018.

[7] Nick Dexter, Clayton G. Webster, and Guannan Zhang. Explicit cost bounds of stochastic Galerkin approximations for parameterized PDEs with random coefficients. *Computers and Mathematics with Applications*, 71(11):2231–2256, 2016.

[8] A. Doostan and H. Owhadi. A non-adapted sparse approximation of PDEs with stochastic inputs. *Journal of Computational Physics*, 230:3015–3034, 2011.

[9] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis. Birkhäuser, 2013.

[10] M. Gunzburger, C. G. Webster, and G. Zhang. Stochastic finite element methods for partial differential equations with random input data. *Acta Numerica*, 23:521–650, 2014.

[11] E. Hale, W. Yin, and Y. Zhang. Fixed-point continuation for ℓ_1 -minimization: methodology and convergence. *SIAM J. Optim.*, 19(3):1107–1130, 2008.

[12] J. Hampton and A. Doostan. Compressive sampling of polynomial chaos expansions: Convergence analysis and sampling strategies. *Journal of Computational Physics*, 280:363–386, 2015.

[13] L. Mathelin and K. Gallivan. A compressed sensing approach for partial differential equations with random input data. *Commun. Comput. Phys.*, 12:919–954, 2012.

[14] F. Nobile, R. Tempone, and C. Webster. A sparse grid stochastic collocation method for elliptic partial differential equations with random input data. *SIAM J. Numer. Anal.*, 46:2309–2345, 2008.

[15] Ji Peng, Jerrad Hampton, and Alireza Doostan. A weighted ℓ_1 -minimization approach for sparse polynomial chaos expansions. *Journal of Computational Physics*, 267:92–111, 2014.

[16] Ji Peng, Jerrad Hampton, and Alireza Doostan. On polynomial chaos expansion via gradient-enhanced ℓ_1 -minimization. *Journal of Computational Physics*, 310:440 – 458, 2016.

[17] Holger Rauhut and Christoph Schwab. Compressive sensing Petrov-Galerkin approximation of high-dimensional parametric operator equations. *Mathematics of Computation*, 86(304):661–700, 2017.

[18] M. K. Stoyanov and C. G. Webster. A dynamically adaptive sparse grid method for quasi-optimal interpolation of multidimensional functions. *Computers & Mathematics with Applications*, 71(11):2449–2465, 2016.

[19] Hoang Tran, Clayton G. Webster, and Guannan Zhang. Analysis of quasi-optimal polynomial approximations for parameterized PDEs with deterministic and stochastic coefficients. *Numerische Mathematik*, pages 1–43, 2017.

[20] Dongbin Xiu and George Em Karniadakis. The Wiener–Askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computing*, 24(2):619–644, 2002.

[21] X. Yang and G. E. Karniadakis. Reweighted ℓ_1 -minimization method for stochastic elliptic differential equations. *Journal of Computational Physics*, 248:87–108, 2013.