

Compressed Sensing Approaches for Polynomial Approximation of High-Dimensional Functions

Ben Adcock, Simone Brugiapaglia, and Clayton G. Webster

Abstract In recent years, the use of sparse recovery techniques in the approximation of high-dimensional functions has garnered increasing interest. In this work we present a survey of recent progress in this emerging topic. Our main focus is on the computation of polynomial approximations of high-dimensional functions on d -dimensional hypercubes. We show that smooth, multivariate functions possess expansions in orthogonal polynomial bases that are not only approximately sparse but possess a particular type of structured sparsity defined by so-called lower sets. This structure can be exploited via the use of weighted ℓ^1 minimization techniques, and, as we demonstrate, doing so leads to sample complexity estimates that are at most logarithmically dependent on the dimension d . Hence the curse of dimensionality – the bane of high-dimensional approximation – is mitigated to a significant extent. We also discuss several practical issues, including unknown noise (due to truncation or numerical error), and highlight a number of open problems and challenges.

Keywords High-dimensional approximation · Weighted ℓ^1 minimization · Orthogonal polynomials · Lower sets

1 Introduction

The approximation of high-dimensional functions is a fundamental difficulty in a large number of fields, including neutron, tomographic and magnetic resonance image reconstruction, uncertainty quantification (UQ), optimal control, and parameter identification for engineering and science applications. In addition, this problem

B. Adcock (✉) · S. Brugiapaglia
Simon Fraser University, Burnaby, BC, Canada
e-mail: ben_adcock@sfu.ca; simone_brugiapaglia@sfu.ca

C.G. Webster
University of Tennessee and Oak Ridge National Lab, Oak Ridge, TN, USA
e-mail: webstercg@math.utk.edu; webstercg@ornl.gov

naturally arises in computational solutions to kinetic plasma physics equations, the many-body Schrödinger equation, Dirac and Maxwell equations for molecular electronic structures and nuclear dynamic computations, options pricing equations in mathematical finance, Fokker-Planck and fluid dynamics equations for complex fluids, turbulent flow, quantum dynamics, molecular life sciences, and nonlocal mechanics. The subject of intensive research over the last half-century, high-dimensional approximation is made challenging by the *curse of dimensionality*, a phrase coined by Bellman [7]. Loosely speaking, this refers to the tendency of naïve approaches to exhibit exponential blow-up in complexity with increasing dimension. Progress is possible, however, by placing restrictions on the class of functions to be approximated, for example, smoothness, anisotropy, sparsity, and compressibility. Well-known algorithms such as sparse grids [14, 54, 55, 69], which are specifically designed to capture this behavior, can mitigate the curse of dimensionality to a substantial extent.

While successful, however, such approaches typically require strong *a priori* knowledge of the functions being approximated, e.g., the parameters of the anisotropic behavior, or costly adaptive implementations to estimate the anisotropy during the approximation process. The efficient approximation of high-dimensional functions in the absence of such knowledge remains a significant challenge.

In this chapter, we consider new methods for high-dimensional approximation based on the techniques of compressed sensing. Compressed sensing is an appealing approach for reconstructing signals from underdetermined systems, i.e., with far smaller number of measurements compared to the signal length [16, 31]. This approach has emerged in the last half a dozen years as an alternative to more classical approximation schemes for high-dimensional functions, with the aim being to overcome some of the limitations mentioned above. Under natural sparsity or compressibility assumptions, it enjoys a significant improvement in sample complexity over traditional methods such as discrete least squares, projection, and interpolation [37, 38]. Our intention in this chapter is to both present an overview of existing work in this area, focusing particularly on the mitigation of the curse of dimensionality, and to highlight existing open problems and challenges.

1.1 Compressed Sensing for High-Dimensional Approximation

Compressed sensing asserts that a vector $\mathbf{x} \in \mathbb{C}^n$ possessing a k -sparse representation in a fixed orthonormal basis can be recovered from a number of suitably chosen measurements m that are linear in k and logarithmic in the ambient dimension n . In practice, recovery can be achieved via a number of different approaches, including convex optimization (ℓ^1 minimization), greedy or thresholding algorithms.

Let $f : D \rightarrow \mathbb{C}$ be a function, where $D \subseteq \mathbb{R}^d$ is a domain in $d \gg 1$ dimensions. In order to apply compressed sensing techniques to approximate f , we must first address the following three questions:

- (i) In which orthonormal system of functions $\{\phi_i\}_{i=1}^n$ does f have an approximately sparse representation?

- (ii) Given suitable assumptions on f (e.g., smoothness) how fast does the best k -term approximation error decay?
- (iii) Given such a system $\{\phi_i\}_{i=1}^n$, what are suitable measurements to take of f ?

The concern of this chapter is the approximation of smooth functions, and as such we will use orthonormal bases consisting of multivariate orthogonal polynomials. In answer to (i) and (ii) in Section 2, we discuss why this choice leads to approximate sparse representations for functions with suitable smoothness and characterize the best k -term approximation error in terms of certain regularity conditions. As we note in Section 2.2, practical examples of such functions include parameter maps of many different types of parametric PDEs.

For sampling, we evaluate f at a set of points $\mathbf{z}_1, \dots, \mathbf{z}_m \in D$. This approach is simple and particularly well-suited in practical problems. In UQ, for example, it is commonly referred to as a *nonintrusive* approach [46] or *stochastic collocation* [52]. More complicated measurement procedures – for instance, *intrusive* procedures such as inner products with respect to a set of basis functions – are often impractical or even infeasible, since, for example, they require computation of high-dimensional integrals. The results presented in Section 3 identify appropriate (random) choices of the sample points $\{\mathbf{z}_i\}_{i=1}^m$ and bound for the number of measurements m under which f can be stably and robustly recovered from the data $\{f(\mathbf{z}_i)\}_{i=1}^m$.

1.2 Structured Sparsity

The approximation of high-dimensional functions using polynomials differs from standard compressed sensing in several key ways. Standard compressed sensing exploits sparsity of the finite vector of coefficients $\mathbf{c} \in \mathbb{C}^n$ of a finite-dimensional signal $\mathbf{x} \in \mathbb{C}^n$. However, polynomial coefficients of smooth functions typically possess more detailed structure than just sparsity. Loosely speaking, coefficients corresponding to low polynomial orders tend to be larger than coefficients corresponding to higher orders. This raises several questions:

- (iv) What is a reasonable *structured sparsity* model for polynomial coefficients of high-dimensional functions?
- (v) How can such structured sparsity be exploited in the reconstruction procedure, and by how much does doing this reduce the number of measurements required?

In Section 2.3 it is shown that high-dimensional functions can be approximated with quasi-optimal rates of convergence by k -term polynomial expansions with coefficients lying in so-called *lower* sets of multi-indices. As we discuss, sparsity in lower sets is a type of structured sparsity, and in Section 3 we show how it can be exploited by replacing the classical ℓ^1 regularizer by a suitable weighted ℓ^1 -norm. Growing weights penalize high-degree polynomial coefficients, and when chosen appropriately, they act to promote lower set structure. In Section 3.2 nonuniform recovery techniques are used to identify a suitable choice of weights. This choice of

weights is then adopted in Section 3.5 to establish quasi-optimal uniform recovery guarantees for compressed sensing of polynomial expansions using weighted ℓ^1 minimization.

The effect of this weighted procedure is a substantially improved recovery guarantee over the case of unweighted ℓ^1 minimization, specifically, a measurement condition that is only logarithmically dependent on the dimension d and polynomially dependent on the sparsity k . Hence the curse of dimensionality is almost completely avoided. As we note in Section 3.3, these polynomial rates of growth in k agree with the best known recovery guarantees for oracle least-squares estimators.

1.3 Dealing with Infinity

Another way in which the approximation of high-dimensional functions differs from standard compressed sensing is that functions typically have infinite (as opposed to finite) expansions in orthogonal polynomial bases. In order to apply compressed sensing techniques, this expansion must be truncated in a suitable way. This leads to the following questions:

- (vi) What is a suitable truncation of the infinite expansion?
- (vii) How does the corresponding truncation error affect the overall reconstruction?

In Section 3 a truncation strategy – corresponding to a hyperbolic cross index set – is proposed based on the lower set structure. The issue of truncation error (question (vii)) presents some technical issues, both theoretical and practical, since this error is usually unknown *a priori*. In Section 3.6 we discuss a means to overcome these issues via a slightly modified optimization problem. Besides doing so, another benefit of the approach developed therein is that it yields approximations to f that also interpolate at the sample points $\{\mathbf{z}_i\}_{i=1}^m$, a desirable property for certain applications. Furthermore, the results given in Section 3.6 also address the robustness of the recovery to unknown errors in the measurements. This is a quite common phenomenon in applications, since function samples are often the result of (inexact) numerical computations.

1.4 Main Results

We now summarize our main results. In order to keep the presentation brief, in this chapter we limit ourselves to functions defined on the unit hypercube $D = (-1, 1)^d$ and consider expansions in orthonormal polynomial bases $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ of Chebyshev or Legendre type. We note in passing, however, that many of our results apply immediately (or extend straightforwardly) to more general systems of functions. See Section 4 for some further discussion.

Let ν be the probability measure under which the basis $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ is orthonormal. Our main result is as follows:

Theorem 1. Let $k \in \mathbb{N}$, $0 < \varepsilon < 1$, $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ be the orthonormal Chebyshev or Legendre basis on $D = (-1, 1)^d$, $\Lambda = \Lambda_k^{\text{HC}}$ be the hyperbolic cross of index k and define weights $\mathbf{u} = (u_i)_{i \in \mathbb{N}_0^d}$, where $u_i = \|\phi_i\|_{L^\infty}$. Suppose that

$$m \gtrsim k^\gamma (\log^2(k) \min\{d + \log(k), \log(2d) \log(k)\} + \log(k) \log(\log(k)/\varepsilon)),$$

where $\gamma = \frac{\log(3)}{\log(2)}$ or $\gamma = 2$ for Chebyshev or Legendre polynomials, respectively, and draw $\mathbf{z}_1, \dots, \mathbf{z}_m \in D$ independently according to ν . Then with probability at least $1 - \varepsilon$ the following holds. For any $f \in L_\nu^2(D) \cap L^\infty(D)$ satisfying

$$\left\| f - \sum_{i \in \Lambda} c_i \phi_i \right\|_{L^\infty} \leq \eta, \quad (1)$$

for some known $\eta \geq 0$, it is possible to compute, via solving a $\ell_{\mathbf{u}}^1$ minimization problem of size $m \times n$ where $n = |\Lambda|$, an approximation \tilde{f} from the samples $\mathbf{y} = (f(\mathbf{z}_j))_{j=1}^m$ that satisfies

$$\|f - \tilde{f}\|_{L_\nu^2} \lesssim \frac{\sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}}}{k^{\gamma/2}} + \eta, \quad \|f - \tilde{f}\|_{L^\infty} \lesssim \sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + k^{\gamma/2} \eta. \quad (2)$$

Here \mathbf{c} are the coefficients of f in the basis $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ and $\sigma_{s,L}(\mathbf{c})_{1,\mathbf{u}}$ is the $\ell_{\mathbf{u}}^1$ -norm error of the best approximation of \mathbf{c} by a vector that is k -sparse and lower.

Note that the condition (1) is strong, since it assumes an *a priori* upper bound on the expansion error is available. Such a condition is unlikely to be met in practice. In Section 3.6 we discuss recovery results for general f without such *a priori* bounds.

1.5 Existing Literature

The first results on compressed sensing with orthogonal polynomials in the one-dimensional setting appeared in [60], based on earlier work in sparse trigonometric expansions [58]. This was extended to the higher-dimensional setting in [72]. Weighted ℓ^1 minimization was studied in [61], and recovery guarantees given in terms of so-called weighted sparsity. However, this does not lead straightforwardly to explicit measurement conditions for quasi-best k -term approximation. The works [1, 22] introduced new guarantees for weighted ℓ^1 minimization of nonuniform and uniform types, respectively, leading to optimal sample complexity estimates for recovering high-dimensional functions using k -term approximations in lower sets. Theorem 1 is based on results in [22]. Relevant approaches to compressed sensing in infinite dimensions have also been considered in [1, 2, 4, 11, 13, 66]

Applications of compressed sensing to UQ, specifically the computation of polynomial chaos expansions of parametric PDEs, can be found in [10, 32, 47, 56, 59, 73]

and references therein. Throughout this chapter we use random sampling from the orthogonality measure of the polynomial basis. We do this for its simplicity, and the theoretical optimality of the recovery guarantees in terms of the dimension d . Other strategies, which typically seek a smaller error or lower polynomial factor of k in the sample complexity, have been considered in [39, 41, 44, 52, 53, 64, 71]. Working toward a similar end, various approaches have also been considered to learn a better sparsity basis [43, 74] or to use additional gradient samples [57]. In this chapter, we focus on fixed bases of Chebyshev or Legendre polynomials in the unit cube. For results in \mathbb{R}^d using Hermite polynomials, see [39, 41, 53].

In some scenarios, a suitable lower set may be known in advance or be computed via an adaptive search. In this case, least-squares methods may be suitable. A series of works have studied the sample complexity of such approaches in the context of high-dimensional polynomial approximation [19, 24, 28, 29, 40, 48–51, 53]. We review a number of these results in Section 3.3.

2 Sparse Polynomial Approximation of High-Dimensional Functions

2.1 Setup and Notation

We first require some notation. For the remainder of this chapter, $D = (-1, 1)^d$ will be the d -dimensional unit cube. The vector $\mathbf{z} = (z_1, \dots, z_d)$ will denote the variable in D , and $\mathbf{i} = (i_1, \dots, i_d) \in \mathbb{N}_0^d$ will be a multi-index. Let $\nu^{(1)}, \dots, \nu^{(d)}$ be probability measures on the unit interval $(-1, 1)$. We consider the tensor product probability measure ν on D given by $\nu = \nu^{(1)} \otimes \dots \otimes \nu^{(d)}$. Let $\{\phi_i^{(k)}\}_{i=0}^\infty$ be an orthonormal polynomial basis of $L_{\nu^{(k)}}^2(-1, 1)$ and define the corresponding tensor product orthonormal basis $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ of $L_\nu^2(D)$ by

$$\phi_{\mathbf{i}} = \phi_{i_1}^{(1)} \otimes \dots \otimes \phi_{i_d}^{(d)}, \quad \mathbf{i} = (i_1, \dots, i_d) \in \mathbb{N}_0^d.$$

We let $\|\cdot\|_{L_\nu^2}$ and $\langle \cdot, \cdot \rangle_{L_\nu^2}$ denote the norm and inner product on $L_\nu^2(D)$, respectively.

Let $f \in L_\nu^2(D) \cap L^\infty(D)$ be the function to be approximated, and write

$$f = \sum_{\mathbf{i} \in \mathbb{N}_0^d} c_i \phi_{\mathbf{i}}, \tag{3}$$

where $c_i = \langle f, \phi_{\mathbf{i}} \rangle_{L_\nu^2}$ are the coefficients of f in the basis $\{\phi_{\mathbf{i}}\}_{i \in \mathbb{N}_0^d}$. We define

$$\mathbf{c} = (c_i)_{i \in \mathbb{N}_0^d} \in \ell^2(\mathbb{N}_0^d),$$

to be the infinite vector of coefficients in this basis.

Example 1. Our main example will be Chebyshev or Legendre polynomials. In one dimension, these are orthogonal polynomials with respect to the weight functions

$$dv = \frac{1}{2} dz \quad (\text{Legendre}), \quad dv = \frac{1}{\pi \sqrt{1-z^2}} dz \quad (\text{Chebyshev}),$$

respectively. For simplicity, we will consider only tensor products of the same types of polynomials in each coordinate. The corresponding tensor product measures on D are consequently defined as:

$$dv = 2^{-d} dz \quad (\text{Legendre}), \quad dv = \prod_{j=1}^d \frac{1}{\pi \sqrt{1-z_j^2}} dz \quad (\text{Chebyshev}).$$

We note also that many of the results presented below extend to more general families of orthogonal polynomials, e.g., Jacobi polynomials (see Remark 5).

As discussed in Section 1.3, it is necessary to truncate the infinite expansion (3) to a finite one. Throughout, we let $\Lambda \subset \mathbb{N}_0^d$ be a subset of size $|\Lambda| = n$, and define the truncated expansion

$$f_\Lambda = \sum_{i \in \Lambda} c_i \phi_i.$$

We write \mathbf{c}_Λ for the finite vector of coefficients with multi-indices in Λ . Whenever necessary, we will assume an ordering i_1, \dots, i_n of the multi-indices in Λ , so that

$$f_\Lambda = \sum_{j=1}^n c_{i_j} \phi_{i_j}, \quad \mathbf{c}_\Lambda = (c_{i_j})_{j=1}^n \in \mathbb{C}^n.$$

We will adopt the usual convention and view \mathbf{c}_Λ interchangeably as a vector in \mathbb{C}^n and as an element of $\ell^2(\mathbb{N}_0^d)$ whose entries corresponding to indices $i \notin \Lambda$ are zero.

2.2 Regularity and Best k -Term Approximation

In the high-dimensional setting, we assume the regularity of f is such that the complex continuation of f , represented as the map $f : \mathbb{C}^d \rightarrow \mathbb{C}$, is a holomorphic function on \mathbb{C}^d . In addition, for $1 \leq k \leq n$, we let

$$\Sigma_k = \{\mathbf{c} \in \ell^2(\mathbb{N}_0^d) : |\text{supp}(\mathbf{c})| \leq k\},$$

be the set of k -sparse vectors, and

$$\sigma_k(\mathbf{c})_1 = \inf_{\mathbf{d} \in \Sigma_k} \|\mathbf{c} - \mathbf{d}\|_1,$$

be the error of the best k -term approximation of \mathbf{c} , measured in the ℓ^1 norm.

Recently, for smooth functions as described above, sparse recovery of the polynomial expansion (3) with the use of compressed sensing has shown tremendous promise. However, this approach requires a small uniform bound of the underlying basis, given by

$$\Theta = \sup_{i \in \Lambda} \|\phi_i\|_{L^\infty(D)},$$

as the sample complexity m required to recover the best k -term approximation (up to a multiplicative constant) scales with the following bound (see, e.g., [35])

$$m \gtrsim \Theta^2 k \times \log \text{ factors}. \quad (4)$$

This poses a challenge for many multivariate polynomial approximation strategies as Θ is prohibitively large in high dimensions. In particular, for d -dimensional problems, $\Theta = 2^{d/2}$ for Chebyshev systems and so-called preconditioned Legendre systems [60]. Moreover, when using the standard Legendre expansion, the theoretical number of samples can exceed the cardinality of the polynomial subspace, unless the subspace a priori excludes all terms of high total order (see, e.g., [41, 72]). Therefore, the advantages of sparse polynomial recovery methods, coming from reduced sample complexity, are eventually overcome by the curse of dimensionality, in that such techniques require at least as many samples as traditional sparse interpolation techniques in high dimensions [37, 54, 55]. Nevertheless, in the next section we describe a common characteristic of the polynomial space spanned by the best k -terms, that we will exploit to overcome the curse of dimensionality in the sample complexity bound (4). As such, our work also provides a fair comparison with existing numerical polynomial approaches in high dimensions [6, 18–20, 65].

2.3 Lower Sets and Structured Sparsity

In many engineering and science applications, the target functions, despite being high-dimensional, are smooth and often characterized by a rapidly decaying polynomial expansion, whose most important coefficients are of low order [21, 25, 27, 42]. In such situations, the quest for finding the approximation containing the largest k terms can be restricted to polynomial spaces associated with *lower* (also known as *downward closed* or *monotone*) sets. These are defined as follows:

Definition 1. A set $S \subseteq \mathbb{N}_0^d$ is lower if, whenever $\mathbf{i} = (i_1, \dots, i_d) \in S$ and $\mathbf{i}' = (i'_1, \dots, i'_d) \in \mathbb{N}_0^d$ satisfies $i'_j \leq i_j$ for all $j = 1, \dots, d$, then $\mathbf{i}' \in S$.

The practicality of downward closed sets is mainly computational, and has been demonstrated in different approaches such as quasi-optimal strategies, Taylor expansion, interpolation methods, and discrete least squares (see [6, 18–22, 26, 27, 49, 51, 62, 65] and references therein). For instance, in the context of parametric PDEs, it was shown in [21] that for a large class of smooth differential operators, with a certain type of anisotropic dependence on \mathbf{z} , the solution map $\mathbf{z} \mapsto f(\mathbf{z})$ can be approximated by its best k -term expansions associated with index sets of cardinality k , resulting in algebraic rates $k^{-\alpha}$, $\alpha > 0$ in the uniform and/or mean average sense. The same rates are preserved with index sets that are lower. In addition, such lower sets of cardinality k also enable the equivalence property $\|\cdot\|_{L^2_\nu(D)} \leq \|\cdot\|_{L^\infty} \leq k^\gamma \|\cdot\|_{L^2_\nu(D)}$ in arbitrary dimensions d with, e.g., $\gamma = 2$ for the uniform measure and $\gamma = \frac{\log 3}{\log 2}$ for Chebyshev measure.

Rather than best k -term approximation, we now consider best k -term approximation in a lower set. Hence, we replace Σ_k with

$$\Sigma_{k,L} = \{\mathbf{c} \in \ell^2(\mathbb{N}_0^d) : |\text{supp}(\mathbf{c})| \leq k, \text{supp}(\mathbf{c}) \text{ is lower}\},$$

and $\sigma_k(\mathbf{c})_1$ with the quantity

$$\sigma_{k,L}(\mathbf{c})_{1,\mathbf{w}} = \inf_{\mathbf{d} \in \Sigma_{k,L}} \|\mathbf{c} - \mathbf{d}\|_{1,\mathbf{w}}. \quad (5)$$

Here $\mathbf{w} = (w_i)_{i \in \mathbb{N}_0^d}$ is a sequence of positive weights and $\|\mathbf{c}\|_{1,\mathbf{w}} = \sum_{i \in \mathbb{N}_0^d} w_i |c_i|$ is the norm on $\ell^1_{\mathbf{w}}(\mathbb{N}_0^d)$.

Remark 1. Sparsity in lower sets is an example of a so-called *structured sparsity* model. Specifically, $\Sigma_{k,L}$ is the subset of Σ_k corresponding to the union of all k -dimensional subspaces defined by lower sets:

$$\Sigma_{k,L} \equiv \bigcup_{\substack{|S|=k \\ S \text{ lower}}} \{\mathbf{c} : \text{supp}(\mathbf{c}) \subseteq S\} \subset \bigcup_{|S|=k} \{\mathbf{c} : \text{supp}(\mathbf{c}) \subseteq S\} \equiv \Sigma_k.$$

Structured sparsity models have been studied extensively in compressed sensing (see, e.g., [5, 9, 30, 33, 66] and references therein). There are a variety of general approaches for exploiting such structure, including greedy and iterative methods (see, for example, [5]) and convex relaxations [66]. A difficulty with lower set sparsity is that projections onto $\Sigma_{k,L}$ cannot be easily computed [22], unlike the case of Σ_k . Therefore, in this chapter we shall opt for a different approach based on $\ell^1_{\mathbf{w}}$ minimization with suitably chosen weights \mathbf{w} . See Section 4 for some further discussion.

3 Compressed Sensing for Multivariate Polynomial Approximation

Having introduced tensor orthogonal polynomials as a good basis for obtaining (structured) sparse representation of smooth, multivariate functions, we now turn our attention to computing quasi-optimal approximations of such a function f from the measurements $\{f(\mathbf{z}_i)\}_{i=1}^m$.

It is first necessary to choose the sampling points $\mathbf{z}_1, \dots, \mathbf{z}_m$. From now on, following an approach that has become standard in compressed sensing [35], we shall assume that these points are drawn randomly and independently according to the probability measure ν . We remark in passing that this may not be the best choice in practice. However, such an approach yields recovery guarantees with measurement conditions that are essentially independent of d , thus mitigating the curse of dimensionality. In Section 4 we briefly discuss other strategies for choosing these points which may convey some practical advantages.

3.1 Exploiting Lower Set-Structured Sparsity

Let $\mathbf{c} \in \ell^2(\mathbb{N}_0^d)$ be the infinite vector of coefficients of a function $f \in L_\nu^2(D)$. Suppose that $\Lambda \subset \mathbb{N}_0^d$, $|\Lambda| = n$ and notice that

$$\mathbf{y} = A\mathbf{c}_\Lambda + \mathbf{e}_\Lambda, \quad (6)$$

where $\mathbf{y} \in \mathbb{C}^m$ and $A \in \mathbb{C}^{m \times n}$ are the finite vector and matrix given by

$$\mathbf{y} = \frac{1}{\sqrt{m}} (f(\mathbf{z}_j))_{j=1}^m, \quad A = \frac{1}{\sqrt{m}} (\phi_{i_k}(\mathbf{z}_j))_{j,k=1}^{m,n}, \quad (7)$$

respectively, and

$$\mathbf{e}_\Lambda = \frac{1}{\sqrt{m}} (f(\mathbf{z}_j) - f_\Lambda(\mathbf{z}_j))_{j=1}^m = \frac{1}{\sqrt{m}} \left(\sum_{i \notin \Lambda} c_i \phi_i(\mathbf{z}_j) \right)_{j=1}^m, \quad (8)$$

is the vector of remainder terms corresponding to the coefficients c_i with indices outside Λ . Our aim is to approximate \mathbf{c} up to an error depending on $\sigma_{k,L}(\mathbf{c})_{1,w}$, i.e., its best k -term approximation in a lower set (see (5)). In order for this to be possible, it is necessary to choose Λ so that it contains all lower sets of cardinality k . A straightforward choice is to make Λ exactly equal to the union of all such sets, which transpires to be precisely the hyperbolic cross index set with index k . That is,

$$\bigcup_{\substack{|S|=k \\ S \text{ lower}}} S = \left\{ \mathbf{i} = (i_1, \dots, i_d) \in \mathbb{N}_0^d : \prod_{j=1}^d (i_j + 1) \leq k \right\} = \Lambda_k^{\text{HC}}. \quad (9)$$

It is interesting to note that this union is a finite set, due to the lower set assumption. Had one not enforced this additional property, the union would be infinite and equal to the whole space \mathbb{N}_0^d .

We shall assume that $\Lambda = \Lambda_k^{\text{HC}}$ from now on. For later results, it will be useful to know the cardinality of this set. While an exact formula in terms of k and d is unknown, there are a variety of different upper and lower bounds. In particular, we shall make use of the following result:

$$n = |\Lambda_k^{\text{HC}}| \leq \min \{2k^3 4^d, e^2 k^{2+\log_2(d)}\}. \quad (10)$$

See [17, Thm. 3.7] and [45, Thm. 4.9], respectively.

With this in hand, we now wish to obtain a solution $\hat{\mathbf{c}}_\Lambda$ of (6) which approximates \mathbf{c}_Λ , and therefore \mathbf{c} due to the choice of Λ , up to an error determined by its best approximation in a lower set of size k . We shall do this by weighted ℓ^1 minimization. Let $\mathbf{w} = (w_i)_{i \in \Lambda}$ be a vector of positive weights and consider the problem

$$\min_{\mathbf{d} \in \mathbb{C}^n} \|\mathbf{d}\|_{1, \mathbf{w}} \text{ s.t. } \|\mathbf{y} - \mathbf{A}\mathbf{d}\|_2 \leq \eta, \quad (11)$$

where $\|\mathbf{d}\|_{1, \mathbf{w}} = \sum_{j=1}^n w_j |d_j|$ is the weighted ℓ^1 -norm and $\eta \geq 0$ is a parameter that will be chosen later. Since the weights \mathbf{w} are positive we shall without loss of generality assume that

$$w_i \geq 1, \quad \forall i.$$

Our choice of these weights is based on the desire to exploit the lower set structure. Indeed, since lower sets inherently penalize higher indices, it is reasonable (and will turn out to be the case) that appropriate choices of increasing weights will promote this type of structure.

For simplicity, we shall assume for the moment that η is chosen so that

$$\eta \geq \|\mathbf{e}_\Lambda\|_2. \quad (12)$$

In particular, this implies that the exact vector \mathbf{c}_Λ is a feasible point of the problem (11). As was already mentioned in Section 1.4, this assumption is a strong one and is unreasonable for practical scenarios where good *a priori* estimates on the expansion error $f - f_\Lambda$ are hard to obtain. In Section 3.6 we address the removal of this condition.

3.2 Choosing the Optimization Weights: Nonuniform Recovery

Our first task is to determine a good choice of optimization weights. For this, techniques from nonuniform recovery¹ are particularly useful.

At this stage it is convenient to define the following. First, for a vector of weights \mathbf{w} and a subset S we let

$$|S|_{\mathbf{w}} = \sum_{i \in S} w_i^2, \quad (13)$$

be the *weighted* cardinality of S . Second, for the orthonormal basis $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ we define the *intrinsic* weights $\mathbf{u} = (u_i)_i$ as

$$u_i = \|\phi_i\|_{L^\infty}. \quad (14)$$

Note that $u_i = \|\phi_i\|_{L^\infty} \geq \|\phi_i\|_{L_v^2} = 1$ since ν is a probability measure. With this in hand, we now have the following result (see [1, Thm. 6.1]):

Theorem 2. *Let $\Lambda \subset \mathbb{N}_0^d$ with $|\Lambda| = n$, $0 < \varepsilon < e^{-1}$, $\eta \geq 0$, $\mathbf{w} = (w_i)_{i \in \Lambda}$ be a set of weights, $\mathbf{c} \in \ell^2(\mathbb{N}_0^d)$ and $S \subseteq \Lambda$, $S \neq \emptyset$, be any fixed set. Draw $\mathbf{z}_1, \dots, \mathbf{z}_m$ independently according to the measure ν , let A , \mathbf{y} and \mathbf{e}_Λ be as in (7) and (8), respectively, and suppose that η satisfies (12). Then, with probability at least $1 - \varepsilon$, any minimizer $\hat{\mathbf{c}}_\Lambda$ of (11) satisfies*

$$\|\mathbf{c} - \hat{\mathbf{c}}_\Lambda\|_2 \lesssim \lambda \sqrt{|S|_{\mathbf{w}}} (\eta + \|\mathbf{c} - \mathbf{c}_\Lambda\|_{1,\mathbf{u}}) + \|\mathbf{c} - \mathbf{c}_S\|_{1,\mathbf{w}}, \quad (15)$$

provided

$$m \gtrsim \left(|S|_{\mathbf{u}} + \max_{i \in \Lambda \setminus S} \{u_i^2/w_i^2\} |S|_{\mathbf{w}} \right) L, \quad (16)$$

where $\lambda = 1 + \frac{\sqrt{\log(\varepsilon^{-1})}}{\log(2n\sqrt{|S|_{\mathbf{w}}})}$ and $L = \log(\varepsilon^{-1}) \log(2n\sqrt{|S|_{\mathbf{w}}})$.

Suppose for simplicity that \mathbf{c} were exactly sparse and let $S = \text{supp}(\mathbf{c})$ and $\eta = 0$. Then this result asserts exact recovery of \mathbf{c} , provided the measurement condition (16) holds. Ignoring the log factor L , this condition is determined by

$$\mathcal{M}(S; \mathbf{u}, \mathbf{w}) = |S|_{\mathbf{u}} + \max_{i \in \Lambda \setminus S} \{u_i^2/w_i^2\} |S|_{\mathbf{w}}. \quad (17)$$

¹By nonuniform recovery, we mean results that guarantee recovery of a fixed vector \mathbf{c}_Λ from a single realization of the random matrix A . Conversely, uniform recovery results consider recovery of all sparse (or structured sparse) vectors from a single realization of A . See, for example, [35] for further discussion.

The first term is the weighted cardinality of S with respect to the intrinsic weights \mathbf{u} and is independent of the choice of optimization weights \mathbf{w} . The second term depends on these weights, but the possibly large size of $|S|_{\mathbf{w}}$ is compensated by the factor $\max_{i \in A \setminus S} \{u_i^2/w_i^2\}$.

Seeking to minimize $\mathcal{M}(S; \mathbf{u}, \mathbf{w})$, it is natural to choose the weights \mathbf{w} so that the second term in (17) is equal to the first. This is easily achieved by the choice

$$w_i = u_i, \quad \forall i, \quad (18)$$

with the resulting measurement condition being simply

$$m \gtrsim |S|_{\mathbf{u}} \log(\varepsilon^{-1}) \log(2n\sqrt{|S|_{\mathbf{u}}}). \quad (19)$$

From now on, we primarily consider the weights (18).

Remark 2. Theorem 2 is a nonuniform recovery guarantee for weighted ℓ^1 minimization. Its proof uses the well-known golfing scheme [36], following similar arguments to those given in [4, 15] for unweighted ℓ^1 minimization. Unlike the results in [4, 15], however, it gives a measurement condition in terms of a fixed set S , rather than the sparsity k (or weighted sparsity). In other words, no sparsity (or structured sparsity) model is required at this stage. Such an approach was first pursued in [8] in the context of block sampling in compressed sensing. See also [23].

3.3 Comparison with Oracle Estimators

As noted above, the condition (19) does not require S to be a lower set. In Section 3.4 we shall use this property in order to estimate $|S|_{\mathbf{u}}$ in terms of the sparsity k . First, however, it is informative to compare (19) to the measurement condition of an oracle estimator. Suppose that the set S were known. Then a standard estimator for \mathbf{c} is the least-squares solution

$$\check{\mathbf{c}}_S = (A_S)^\dagger \mathbf{y}, \quad (20)$$

where $A_S \in \mathbb{C}^{m \times |S|}$ is the matrix formed from the columns of A with indices belonging to S and \dagger denotes the pseudoinverse. Stable and robust recovery via this estimator follows if the matrix A_S is well-conditioned. For this, one has the following well-known result:

Proposition 1. *Let $0 < \delta, \varepsilon < 1$, $S \subset \mathbb{N}_0^d$, $|S| = k$ and suppose that m satisfies*

$$m \gtrsim \delta^{-2} |S|_{\mathbf{u}} \log(2k\varepsilon^{-1}). \quad (21)$$

Draw z_1, \dots, z_m independently according to the measure ν and let A be as in (7). Then, with probability at least $1 - \varepsilon$, the matrix A_S satisfies

$$\|(A_S)^* A_S - I\|_2 \leq \delta,$$

where $I \in \mathbb{C}^{k \times k}$ is the identity matrix and $\|\cdot\|_2$ is the matrix 2-norm.

See, for example, [1, Lem. 8.2]. Besides the log factor, (21) is the same sufficient condition as (19). Thus the weighted ℓ^1 minimization estimator \hat{c}_A with weights $\mathbf{w} = \mathbf{u}$ requires roughly the same measurement condition as the oracle least-squares estimator. Of course, the former requires no *a priori* knowledge of S .

Remark 3. In fact, one may prove a slightly sharper estimate than (21) where $|S|_{\mathbf{u}}$ is replaced by the quantity

$$\sup_{z \in D} \sum_{i \in S} |\phi_i(z)|^2. \quad (22)$$

See, for example, [24]. Note that $\sum_{i \in S} |\phi_i(z)|^2$ is the so-called Christoffel function of the subspace spanned by the functions $\{\phi_i\}_{i \in S}$. However, (22) coincides with $|S|_{\mathbf{u}}$ whenever the polynomials ϕ_i achieve their absolute maxima at the same point in D . This is the case for any Jacobi polynomials whenever the parameters satisfy $\max\{\alpha, \beta\} \geq -1/2$ [63, Thm. 7.32.1]; in particular, Legendre and Chebyshev polynomials (see Example 1), and tensor products thereof.

3.4 Sample Complexity for Lower Sets

The measurement condition (19) determines the sample complexity in terms of the weighted cardinality $|S|_{\mathbf{u}}$ of the set S . When a structured sparsity model is applied to S – in particular, lower set sparsity – one may derive estimates for $|S|_{\mathbf{u}}$ in terms of just the cardinality $k = |S|$ and the dimension d .

Lemma 1. *Let $2 \leq k \leq 2^{d+1}$. If $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ is the tensor Chebyshev basis then*

$$k^{\log(3)/\log(2)}/3 \leq \max \{ |S|_{\mathbf{u}} : S \subset \mathbb{N}_0^d, |S| \leq k, S \text{ lower} \} \leq k^{\log(3)/\log(2)},$$

where $|S|_{\mathbf{u}}$ and \mathbf{u} are as in (13) and (14), respectively. If $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ is the tensor Legendre basis then

$$k^2/4 \leq \max \{ |S|_{\mathbf{u}} : S \subset \mathbb{N}_0^d, |S| \leq k, S \text{ lower} \} \leq k^2.$$

Moreover, the upper estimates hold for all $k \geq 2$.

See [19, Lem. 3.7]. With this in hand, we now have the following result:

Theorem 3. Consider the setup in Theorem 2 with $k \geq 2$, $\Lambda = \Lambda_k^{\text{HC}}$ the hyperbolic cross (9), weights $\mathbf{w} = \mathbf{u}$, and $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ the tensor Legendre or Chebyshev basis. Then any minimizer $\hat{\mathbf{c}}_\Lambda$ of (11) with weights $\mathbf{w} = \mathbf{u}$ satisfies

$$\|\mathbf{c} - \hat{\mathbf{c}}_\Lambda\|_2 \lesssim \lambda k^{\gamma/2} (\eta + \|\mathbf{c} - \mathbf{c}_\Lambda\|_{1,\mathbf{u}}) + \sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}},$$

with probability at least $1 - \varepsilon$, provided

$$m \gtrsim k^\gamma \log(\varepsilon^{-1}) \min \{d + \log(k), \log(2d) \log(k)\},$$

where $\lambda = 1 + \frac{\sqrt{\log(\varepsilon^{-1})}}{\log(k)}$ and where $\gamma = \log(3)/\log(2)$ or $\gamma = 2$ in the Chebyshev or Legendre case, respectively.

Proof. Let $S \subset \mathbb{N}_0^d$, $|S| \leq k$ be a lower set such that $\|\mathbf{c} - \mathbf{c}_S\|_{1,\mathbf{u}} = \sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}}$. By Lemma 1 we have $|S|_{\mathbf{u}} \leq k^\gamma$. We now apply Theorem 2 with $\mathbf{w} = \mathbf{u}$, and use this result and the bound (10) for $n = |\Lambda_k^{\text{HC}}|$. \square

Remark 4. It is worth noting that the lower set assumption drastically reduces the sample complexity. Indeed, for the case of Chebyshev polynomials one has

$$\max \{|S|_{\mathbf{u}} : S \subset \mathbb{N}_0^d, |S| \leq k\} = 2^d k.$$

In other words, in the absence of the lower set condition, one can potentially suffer exponential blow-up with dimension d . Note that this result follows straightforwardly from the explicit expression for the weights \mathbf{u} in this case: namely,

$$u_i = 2^{\|\mathbf{i}\|_0/2}, \quad (23)$$

where $\|\mathbf{i}\|_0 = |\{j : i_j \neq 0\}|$ for $\mathbf{i} = (i_1, \dots, i_d) \in \mathbb{N}_0^d$ (see, e.g., [1]). On the other hand, for Legendre polynomials the corresponding quantity is infinite, since in this case the weights

$$u_i = \prod_{j=1}^d \sqrt{2i_j + 1}, \quad (24)$$

are unbounded. Moreover, even if S is constrained to lie in the hyperbolic cross $\Lambda = \Lambda_k^{\text{HC}}$, one still has a worst-case estimate that is polynomially large in k [22].

Remark 5. We have considered only tensor Legendre and Chebyshev polynomial bases. However, Theorem 3 readily extends to other types of orthogonal polynomials. All that is required is an upper bound for

$$\max \{|S|_{\mathbf{u}} : S \subset \mathbb{N}_0^d, |S| \leq k, S \text{ lower}\},$$

in terms of the sparsity k . For example, suppose that $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ is the tensor ultraspherical polynomial basis, corresponding to the measure

$$d\nu = (c_\alpha)^d \prod_{j=1}^d (1 - z_j^2)^\alpha dz, \quad c_\alpha = \left(\int_{-1}^1 (1 - z^2)^\alpha dz \right)^{-1}.$$

If the parameter α satisfies $2\alpha + 1 \in \mathbb{N}$ then [49, Thm. 8] gives that

$$\max \{|S|_u : S \subset \mathbb{N}_0^d, |S| \leq k, S \text{ lower}\} \leq k^{2\alpha+2}.$$

This result includes the Legendre case ($\alpha = 0$) given in Lemma 1, as well as the case of Chebyshev polynomials of the second kind ($\alpha = 1/2$). A similar result also holds for tensor Jacobi polynomials for parameters $\alpha, \beta \in \mathbb{N}_0$ (see [49, Thm. 9]).

3.5 Quasi-Optimal Approximation: Uniform Recovery

As is typical of a nonuniform recovery guarantee, the error bound in Theorem 3 has the limitation that it relates the ℓ^2 -norm of the error with the best k -term, lower approximation error in the ℓ_u^1 -norm. To obtain stronger estimates we now consider uniform recovery techniques.

We first require an extension of the standard restricted isometry property (RIP) to the case of sparsity in lower sets. To this end, for $k \in \mathbb{N}$ we now define the quantity

$$s(k) = \max \{|S|_u : S \subset \mathbb{N}_0^d, |S| \leq k, S \text{ lower}\}. \tag{25}$$

The following extension of the RIP was introduced in [22]:

Definition 2. A matrix $A \in \mathbb{C}^{m \times n}$ has the lower restricted isometry property (lower RIP) of order k if there exists a $0 < \delta < 1$ such that

$$(1 - \delta) \|c\|_2^2 \leq \|Ac\|_2^2 \leq (1 + \delta) \|c\|_2^2, \quad \forall c \in \mathbb{C}^n, |\text{supp}(c)|_u \leq s(k).$$

If $\delta = \delta_{k,L}$ is the smallest constant such that this holds, then $\delta_{k,L}$ is the k^{th} lower restricted isometry constant (lower RIC) of A .

We shall use the lower RIP to establish stable and robust recovery. For this, we first note that the lower RIP implies a suitable version of the robust null space property (see [22, Prop. 4.4]):

Lemma 2. Let $k \geq 2$ and $A \in \mathbb{C}^{m \times n}$ satisfy the lower RIP of order αk with constant

$$\delta = \delta_{\alpha k, L} < 1/5,$$

where $\alpha = 2$ if the weights \mathbf{u} arise from the tensor Legendre basis and $\alpha = 3$ if the weights arise from the tensor Chebyshev basis. Then for any $S \subseteq \Lambda_s^{\text{HC}}$ with $|S|_{\mathbf{u}} \leq s(k)$ and any $\mathbf{d} \in \mathbb{C}^n$,

$$\|\mathbf{d}_S\|_2 \leq \frac{\rho}{\sqrt{s(k)}} \|\mathbf{d}_{S^c}\|_{1,\mathbf{u}} + \tau \|\mathbf{A}\mathbf{d}\|_2,$$

where $\rho = \frac{4\delta}{1-\delta} < 1$ and $\tau = \frac{\sqrt{1+\delta}}{1-\delta}$.

With this in hand, we now establish conditions under which the lower RIP holds for matrices A defined in (7). The following result was shown in [22]:

Theorem 4. Fix $0 < \varepsilon < 1$, $0 < \delta < 1/13$, let $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ be as in Section 2.1 and \mathbf{u} be as in (14) and suppose that

$$m \gtrsim \frac{s(k)}{\delta^2} L,$$

where $s(k)$ is as in (25) and

$$L = \log\left(\frac{s(k)}{\delta^2}\right) \left(\frac{1}{\delta^4} \log\left(2 \frac{s(k)}{\delta^2} \log\left(\frac{s(k)}{\delta^2}\right)\right) \log(2n) + \frac{1}{\delta} \log\left(\frac{1}{\gamma\delta} \log\left(\frac{k(s)}{\delta^2}\right)\right)\right).$$

Draw $\mathbf{z}_1, \dots, \mathbf{z}_m$ independently according to ν and let $A \in \mathbb{C}^{m \times n}$ be as in (7). Then with probability at least $1 - \varepsilon$, the matrix A satisfies the lower RIP of order k with constant $\delta_{k,L} \leq 13\delta$.

Combining this with the previous lemma now gives the following uniform recovery guarantee:

Theorem 5. Let $0 < \varepsilon < 1$, $k \geq 2$ and

$$m \asymp k^\gamma L, \tag{26}$$

where $\gamma = \log(3)/\log(2)$ or $\gamma = 2$ in the tensor Chebyshev or tensor Legendre cases, respectively, and

$$L = (\log^2(k) \min\{d + \log(k), \log(2d) \log(k)\} + \log(k) \log(\log(k)/\varepsilon)). \tag{27}$$

Let $\Lambda = \Lambda_k^{\text{HC}}$ be the hyperbolic cross index set, $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ be the tensor Legendre or Chebyshev polynomial basis and draw $\mathbf{z}_1, \dots, \mathbf{z}_m$ independently according to the corresponding measure ν . Then with probability at least $1 - \varepsilon$ the following holds. For any $f \in L^2(D) \cap L^\infty(D)$ the approximation

$$\tilde{f} = \sum_{i \in \Lambda} \hat{c}_i \phi_i,$$

where $\hat{\mathbf{c}}_\Lambda = (\hat{c}_i)_{i \in \Lambda}$ is a solution of (11) with A , \mathbf{y} and η given by (7) and (12), respectively, and weights $\mathbf{w} = \mathbf{u}$, satisfies

$$\|f - \tilde{f}\|_{L^\infty} \leq \|\mathbf{c} - \hat{\mathbf{c}}_\Lambda\|_{1,\mathbf{u}} \lesssim \sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + k^{\gamma/2}\eta, \quad (28)$$

and

$$\|f - \tilde{f}\|_{L^2_\nu} = \|\mathbf{c} - \hat{\mathbf{c}}_\Lambda\|_2 \lesssim \frac{\sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}}}{k^{\gamma/2}} + \eta, \quad (29)$$

where $\mathbf{c} \in \ell^2(\mathbb{N}_0^d)$ are the coefficients of f in the basis $\{\phi_i\}_{i \in \mathbb{N}_0^d}$.

Proof. Let $\alpha = 2$ or $\alpha = 3$ in the Legendre or Chebyshev case, respectively. Condition (26), Lemma 1 and Theorem 4 imply that the matrix A satisfies the lower RIP of order αk with constant $\delta_{\alpha k,L} \leq 1/6 < 1/5$. Now let S be a lower set of cardinality $|S| = k$ such that

$$\|\mathbf{c} - \mathbf{c}_S\|_{1,\mathbf{u}} = \sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}}, \quad (30)$$

set $\mathbf{d} = \mathbf{c}_\Lambda - \hat{\mathbf{c}}_\Lambda$ and $T = \Lambda \setminus S$. Note that

$$\begin{aligned} \|\mathbf{d}_T\|_{1,\mathbf{u}} &\leq \|\mathbf{c}_T\|_{1,\mathbf{u}} + \|\hat{\mathbf{c}}_T\|_{1,\mathbf{u}} \\ &= 2\|\mathbf{c}_T\|_{1,\mathbf{u}} + \|\mathbf{c}_S\|_{1,\mathbf{u}} + \|\hat{\mathbf{c}}_T\|_{1,\mathbf{u}} - \|\mathbf{c}_\Lambda\|_{1,\mathbf{u}} \\ &\leq 2\|\mathbf{c}_T\|_{1,\mathbf{u}} + \|\mathbf{d}_S\|_{1,\mathbf{u}} + \|\hat{\mathbf{c}}_\Lambda\|_{1,\mathbf{u}} - \|\mathbf{c}_\Lambda\|_{1,\mathbf{u}} \leq 2\sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + \|\mathbf{d}_S\|_{1,\mathbf{u}}, \end{aligned}$$

since $\hat{\mathbf{c}}_\Lambda$ is a solution of (11) and \mathbf{c}_Λ is feasible for (11) due to the choice of η . By Lemma 2 we have

$$\|\mathbf{d}_T\|_{1,\mathbf{u}} \leq 2\sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + \sqrt{s(k)}\|\mathbf{d}_S\|_2 \leq 2\sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + \rho\|\mathbf{d}_T\|_{1,\mathbf{u}} + \tau\sqrt{s(k)}\|\mathbf{A}\mathbf{d}\|_2,$$

where $\rho \leq 4/5$ and $\tau \leq \sqrt{42}/5$. Therefore

$$\|\mathbf{d}_T\|_{1,\mathbf{u}} \lesssim \sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + \sqrt{s(k)}\|\mathbf{A}\mathbf{d}\|_2 \lesssim \sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + \sqrt{s(k)}\eta,$$

where in the second step we use the fact that $\mathbf{d} = \mathbf{c}_\Lambda - \hat{\mathbf{c}}_\Lambda$ is the difference of two vectors that are both feasible for (11). Using this bound and Lemma 2 again gives

$$\|\mathbf{d}\|_{1,\mathbf{u}} \lesssim \sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + \sqrt{s(k)}\eta, \quad (31)$$

and since $\mathbf{c} - \hat{\mathbf{c}}_\Lambda = \mathbf{d} + \mathbf{c} - \mathbf{c}_\Lambda$, we deduce that

$$\|\mathbf{c} - \hat{\mathbf{c}}_\Lambda\|_{1,\mathbf{u}} \leq \|\mathbf{d}\|_{1,\mathbf{u}} + \|\mathbf{c} - \mathbf{c}_\Lambda\|_{1,\mathbf{u}} \lesssim \sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + \sqrt{s(k)}\eta. \quad (32)$$

Due to the definition of the weights \mathbf{u} , we have $\|f - \tilde{f}\|_{L^\infty} \leq \|\mathbf{c} - \hat{\mathbf{c}}_\Lambda\|_{1,\mathbf{u}}$, and therefore, after noting that $s(k) \lesssim k^\gamma$ (see Lemma 1) we obtain the first estimate (28). For the second estimate let S be such that

$$\|\mathbf{c} - \mathbf{c}_S\|_2 = \min \{ \|\mathbf{c} - \mathbf{d}\|_2 : |\text{supp}(\mathbf{d})|_{\mathbf{u}} \leq s(k) \},$$

and set $T = S^c$. Let $\mathbf{d} = \mathbf{c} - \hat{\mathbf{c}}_\Lambda$ and write $\|\mathbf{d}\|_2 \leq \|\mathbf{d}_S\|_2 + \|\mathbf{d}_T\|_2$. Via a weighted Stechkin estimate [61, Thm. 3.2] we have $\|\mathbf{d}_T\|_2 \leq \frac{1}{\sqrt{s(k) - \|\mathbf{u}\|_\infty}} \|\mathbf{d}\|_{1,\mathbf{u}}$. For tensor Chebyshev and Legendre polynomials, one has $\|\mathbf{u}\|_\infty \leq \frac{3}{4}s(k)$ (see [22, Lem. 4.1]), and therefore $\|\mathbf{d}_T\|_2 \lesssim \frac{1}{\sqrt{s(k)}} \|\mathbf{d}\|_{1,\mathbf{u}}$. We now apply Lemma 2 to deduce that $\|\mathbf{d}\|_2 \lesssim \frac{1}{\sqrt{s(k)}} \|\mathbf{d}\|_{1,\mathbf{u}} + \eta$. Recall that $s(k) \gtrsim k^\gamma$ due to Lemma 1. Hence (32) now gives $\|\mathbf{d}\|_2 \lesssim \frac{\sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}}}{k^{\gamma/2}} + \eta$, as required. \square

For the Legendre and Chebyshev cases, Theorem 5 proves recovery with quasi-optimal k -term rates of approximation subject to the same measurement condition (up to log factors) as the oracle least-squares estimator. In particular, the sample complexity is polynomial in k and at most logarithmic in the dimension d , thus mitigating the curse of dimensionality to a substantial extent. We remark in passing that this result can be extended to general Jacobi polynomials (recall Remark 5). Furthermore, the dependence on d can be removed altogether by considering special classes of lower sets, known as *anchored* sets [29].

3.6 Unknown Errors, Robustness, and Interpolation

A drawback of the main results so far (Theorems 3 and 5) is that they assume the *a priori* bound (12), i.e.

$$\frac{1}{m} \sum_{j=1}^m |f(\mathbf{z}_j) - f_\Lambda(\mathbf{z}_j)|^2 \leq \eta^2, \quad (33)$$

for some known η . Note that this is implied by the slightly stronger condition

$$\|f - f_\Lambda\|_{L^\infty} \leq \eta.$$

Such an η is required in order to formulate the optimization problem (11) to recover f . Moreover, in view of the error bounds in Theorems 3 and 5, one expects a poor estimation of η to yield a larger recovery error. Another drawback of the current approach is that the approximation \tilde{f} does not interpolate f , a property which is sometimes desirable in applications.

We now consider the removal of the condition (12). This follows the work of [3, 12]. To this end, let $\eta \geq 0$ be arbitrary, i.e., (33) need not hold, and consider the minimization problem

$$\min_{\mathbf{d} \in \mathbb{C}^n} \|\mathbf{d}\|_{1,\mathbf{u}} \text{ s.t. } \|\mathbf{y} - A\mathbf{d}\|_2 \leq \eta. \quad (34)$$

If $\hat{\mathbf{c}}_\Lambda = (\hat{c}_i)_{i \in \Lambda}$ is a minimizer of this problem, we define, as before, the corresponding approximation

$$\tilde{f} = \sum_{i \in \Lambda} \hat{c}_i \phi_i.$$

Note that if $\eta = 0$ then \tilde{f} exactly interpolates f at the sample points $\{\mathbf{z}_j\}_{j=1}^m$.

An immediate issue with the minimization problem (34) is that the truncated vector of coefficients \mathbf{c}_Λ is not generally feasible. Indeed, $\mathbf{y} - A\mathbf{c}_\Lambda = \mathbf{e}_\Lambda$, where \mathbf{e}_Λ is as in (8) and is generally nonzero. In fact, is not even guaranteed that the feasibility set of (34) is nonempty. However, this will of course be the case whenever A has full rank m . Under this assumption, one then has the following (see [3]):

Theorem 6. *Let $\varepsilon, k, m, \gamma, \Lambda, \{\phi_i\}_{i \in \mathbb{N}_0^d}$ and $\mathbf{z}_1, \dots, \mathbf{z}_m$ be as in Theorem 5. Then with probability at least $1 - \varepsilon$ the following holds. For any $\eta \geq 0$ and $f \in L^2(D) \cap L^\infty(D)$ the approximation*

$$\tilde{f} = \sum_{i \in \Lambda} \hat{c}_i \phi_i,$$

where $\hat{\mathbf{c}}_\Lambda = (\hat{c}_i)_{i \in \Lambda}$ is a solution of (34) with A and \mathbf{y} given by (7) satisfies

$$\|f - \tilde{f}\|_{L^\infty} \leq \|\mathbf{c} - \hat{\mathbf{c}}_\Lambda\|_{1,\mathbf{u}} \lesssim \sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + k^{\gamma/2} (\eta + \|\mathbf{e}_\Lambda\|_2 + T_u(A, \Lambda, \mathbf{e}_\Lambda, \eta)) \quad (35)$$

and

$$\|f - \tilde{f}\|_{L^2_v} = \|\mathbf{c} - \hat{\mathbf{c}}_\Lambda\|_2 \lesssim \frac{\sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}}}{k^{\gamma/2}} + \eta + \|\mathbf{e}_\Lambda\|_2 + T_u(A, \Lambda, \mathbf{e}_\Lambda, \eta), \quad (36)$$

where $\mathbf{c} \in \ell^2(\mathbb{N}_0^d)$ are the coefficients of f in the basis $\{\phi_i\}_{i \in \mathbb{N}_0^d}$, \mathbf{e}_Λ is as in (8) and

$$T_u(A, \Lambda, \mathbf{e}_\Lambda, \eta) = \min \left\{ \frac{\|\mathbf{d}\|_{1,\mathbf{u}}}{k^{\gamma/2}} : \mathbf{d} \in \mathbb{C}^n, \|A\mathbf{d} - \mathbf{e}_\Lambda\|_2 \leq \eta \right\}. \quad (37)$$

Proof. We follow the steps of the proof of Theorem 5 with some adjustments to take into account the fact that \mathbf{c}_Λ may not be feasible. First, let S be such that (30) holds and set $\mathbf{d} = \mathbf{c}_\Lambda - \hat{\mathbf{c}}_\Lambda$ and $T = \Lambda \setminus S$. Then, arguing in a similar way we see that

$$\begin{aligned} \|\mathbf{d}_T\|_{1,\mathbf{u}} &\leq 2\sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + \|\mathbf{d}_S\|_{1,\mathbf{u}} + \|\hat{\mathbf{c}}_\Lambda\|_{1,\mathbf{u}} - \|\mathbf{c}_\Lambda\|_{1,\mathbf{u}} \\ &\leq 2\sigma_{k,L}(\mathbf{c})_{1,\mathbf{u}} + \|\mathbf{d}_S\|_{1,\mathbf{u}} + \|\mathbf{g} - \mathbf{c}_\Lambda\|_{1,\mathbf{u}}, \end{aligned}$$

where $\mathbf{g} \in \mathbb{C}^n$ is any point in the feasible set of (34). By Lemma 2 we have

$$\|\mathbf{d}_T\|_{1,u} \leq 2\sigma_{k,L}(\mathbf{c})_{1,u} + \rho\|\mathbf{d}_T\|_{1,u} + \tau\sqrt{s(k)}\|\mathbf{A}\mathbf{d}\|_2 + \|\mathbf{g} - \mathbf{c}_\Lambda\|_{1,u}.$$

Notice that $\|\mathbf{A}\mathbf{d}\|_2 = \|\mathbf{y} - \mathbf{e}_\Lambda - \mathbf{A}\hat{\mathbf{c}}_\Lambda\|_2 \leq \|\mathbf{e}_\Lambda\|_2 + \eta$, and therefore

$$\|\mathbf{d}_T\|_{1,u} \lesssim \sigma_{k,L}(\mathbf{c})_{1,u} + \sqrt{s(k)}(\|\mathbf{e}_\Lambda\|_2 + \eta) + \|\mathbf{g} - \mathbf{c}_\Lambda\|_{1,u}.$$

Hence, by similar arguments, it follows that

$$\|\mathbf{c} - \mathbf{c}_\Lambda\|_{1,u} \lesssim \sigma_{k,L}(\mathbf{c})_{1,u} + k^{\gamma/2}(\|\mathbf{e}_\Lambda\|_2 + \eta) + \|\mathbf{g} - \mathbf{c}_\Lambda\|_{1,u}, \quad (38)$$

for any feasible point \mathbf{g} . After analogous arguments, we also deduce the following bound in the ℓ^2 -norm:

$$\|\mathbf{c} - \mathbf{c}_\Lambda\|_2 \lesssim \frac{\sigma_{k,L}(\mathbf{c})_{1,u}}{k^{\gamma/2}} + (\|\mathbf{e}_\Lambda\|_2 + \eta) + k^{-\gamma/2}\|\mathbf{g} - \mathbf{c}_\Lambda\|_{1,u}. \quad (39)$$

To complete the proof, we consider the term $\|\mathbf{g} - \mathbf{c}_\Lambda\|_{1,u}$. Write $\mathbf{g} = \mathbf{c}_\Lambda + \mathbf{g}'$ and notice that \mathbf{g} is feasible if and only if \mathbf{g}' satisfies $\|\mathbf{A}\mathbf{g}' - \mathbf{e}_\Lambda\| \leq \eta$. Since \mathbf{g}' is arbitrary we get the result. \square

The two error bounds (35) and (36) in this theorem are analogous to (28) and (29) in Theorem 5. They remove the condition $\eta \geq \|\mathbf{e}_\Lambda\|_2$ at the expense of an additional term $T_u(\mathbf{A}, \Lambda, \mathbf{e}_\Lambda, \eta)$. We now provide a bound for this term (see [3]):

Theorem 7. *Consider the setup of Theorem 6, and let $T_u(\mathbf{A}, \Lambda, \mathbf{e}_\Lambda, \eta)$ be as in (37). If \mathbf{A} has full rank, then*

$$T_u(\mathbf{A}, \Lambda, \mathbf{e}_\Lambda, \eta) \leq \frac{k^{\alpha/2}\sqrt{L}}{\sigma_{\min}(\sqrt{\frac{m}{n}}\mathbf{A}^*)} \max\{\|\mathbf{e}_\Lambda\|_2 - \eta, 0\}, \quad (40)$$

where L is as in (27) and $\alpha = 1, 2$ in the Chebyshev or Legendre cases, respectively.

Proof. If $\eta \geq \|\mathbf{e}_\Lambda\|_2$ then the result holds trivially. Suppose now that $\eta < \|\mathbf{e}_\Lambda\|_2$. Since $\|\mathbf{e}_\Lambda\|_2 \neq 0$ in this case, we can define $\mathbf{d} = (1 - \eta/\|\mathbf{e}_\Lambda\|_2)\mathbf{A}^\dagger\mathbf{e}_\Lambda$, where \mathbf{A}^\dagger denotes the pseudoinverse. Then \mathbf{d} satisfies $\|\mathbf{A}\mathbf{d} - \mathbf{e}_\Lambda\|_2 = \eta$, and therefore

$$k^{\gamma/2}T_u(\mathbf{A}, \Lambda, \mathbf{e}_\Lambda, \eta) \leq \|\mathbf{d}\|_{1,u} \leq \sqrt{|\Lambda|_u}\|\mathbf{d}\|_2 \leq \frac{\sqrt{|\Lambda|_u}}{\sigma_{\min}(\mathbf{A}^*)}(\|\mathbf{e}_\Lambda\|_2 - \eta).$$

Equation (26) implies that $\sqrt{\frac{m}{k^\gamma}} \lesssim \sqrt{L}$, and hence

$$T_u(\mathbf{A}, \Lambda, \mathbf{e}_\Lambda, \eta) \lesssim \sqrt{\frac{|\Lambda|_{1,u}}{n}} \frac{\sqrt{L}}{\sigma_{\min}(\sqrt{\frac{m}{n}}\mathbf{A}^*)} (\|\mathbf{e}_\Lambda\|_2 - \eta). \quad (41)$$

It remains to estimate $|\Lambda|_{1,u}$. For the Chebyshev case, we apply (23) to give

$$|\Lambda|_{1,\mathbf{u}} = \sum_{i \in \Lambda} 2^{\|i\|_0} \leq \sum_{i \in \Lambda} \prod_{j=1}^d (i_j + 1) \leq k \sum_{i \in \Lambda} 1 = kn$$

where in the penultimate step we used the definition of the hyperbolic cross (9). For the Legendre case, we use (24) to get

$$|\Lambda|_{1,\mathbf{u}} = \sum_{i \in \Lambda} \prod_{j=1}^d (2i_j + 1) \leq \sum_{i \in \Lambda} 2^{\|i\|_0} \prod_{j=1}^d (i_j + 1) \leq k^2 n.$$

This completes the proof. □

The error bound (40) suggests that the effect of removing the condition $\eta \geq \|e_\Lambda\|_2$ is at most a small algebraic factor in k , a log factor and term depending on the minimal singular value of the scaled matrix $\sqrt{\frac{m}{n}}A^*$. We discuss this latter term further in below. Interestingly, this bound suggests that a good estimate of $\|e_\Lambda\|_2$ (when available) can reduce this error term. Indeed, one has $T_u(A, \Lambda, e_\Lambda, \eta) \rightarrow 0$ linearly in $\|e_\Lambda\|_2 - \eta \rightarrow 0^+$. Hence estimation procedures aiming to tune η – for example, cross validation (see Section 3.7) – are expected to yield reduced error over the case $\eta = 0$, for example.

It is beyond the scope of this chapter to provide theoretical bounds on the minimal singular value of the scaled matrix $\sqrt{\frac{m}{n}}A^*$. We refer to [12] for a more comprehensive treatment of such bounds. However, we note that it is reasonable to expect that $\sigma_{\min}(\sqrt{\frac{m}{n}}A^*) \approx 1$ under appropriate conditions on m and n . Indeed:

Lemma 3. *Let $B = \mathbb{E}(\frac{m}{n}AA^*)$, where A is the matrix of Theorem 6. Then the minimal eigenvalue of B is precisely $1 - 1/n$.*

Proof. We have $\mathbb{E}(\frac{m}{n}AA^*)_{j,l} = \mathbb{E}(\frac{1}{n} \sum_{i \in \Lambda} \phi_i(z_j)\phi_i(z_l))$. When $l = j$ this gives $\mathbb{E}(\frac{m}{n}AA^*)_{j,j} = 1$. Conversely, since $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ are orthogonal polynomials one has $\int_D \phi_i(z) \, d\nu = \langle \phi_i, \phi_0 \rangle_{L^2_\nu} = \delta_{i,0}$, and therefore for $l \neq j$ one has $\mathbb{E}(\frac{m}{n}AA^*)_{j,l} = \frac{1}{n} \sum_{i \in \Lambda} (\int_D \phi_i(z) \, d\nu)^2 = \frac{1}{n}$. It is now a straightforward calculation to show that $\lambda_{\min}(B) = 1 - 1/n$. □

Remark 6. Although complete theoretical estimates $T_u(A, \Lambda, e_\Lambda, \eta)$ are outside the scope of this work, it is straightforward to derive a bound that can be computed. Indeed, it follows immediately from (41) that

$$T_u(A, \Lambda, e_\Lambda, \eta) \lesssim Q_u(A) \sqrt{L} \max\{\|e_\Lambda\|_2 - \eta, 0\},$$

where

$$Q_u(A) = \sqrt{\frac{|\Lambda|_{1,\mathbf{u}}}{n} \frac{1}{\sigma_{\min}(\sqrt{\frac{m}{n}}A^*)}}. \tag{42}$$

Hence, up to the log factor, the expected robustness of (34) can be easily checked numerically. See Section 3.7 for some examples of this approach.

Remark 7. For pedagogical reasons, we have assumed the truncation of f to f_Λ is the only source of error e_Λ affecting the measurements \mathbf{y} (recall (8)). There is no reason for this to be the case, and e_Λ may incorporate other errors without changing any of the above results. We note that concrete applications often give rise to other sources of unknown error. For example, in UQ, we usually aim at approximating a function of the form $f(\mathbf{z}) = q(u(\mathbf{z}))$, where $u(\mathbf{z})$ is the solution to a PDE depending on some random coefficients \mathbf{z} and q is a quantity of interest (see, e.g., [32, 73]). In this case, each sample $f(\mathbf{z}_j)$ is typically subject to further sources of inaccuracy, such as the numerical error associated with the PDE solver employed to compute $u(\mathbf{z}_j)$ (e.g., a finite element method) and, possibly, the error committed evaluating q on $u(\mathbf{z}_j)$ (e.g., numerical integration).

Remark 8. Our analysis based on the estimation of the tail error (37) can be compared with the robustness analysis of basis pursuit based on the so-called *quotient property* [35]. However, this analysis is limited to the case of *basis pursuit*, corresponding to the optimization program (34) with $\mathbf{u} = \mathbf{1}$ (i.e., unweighted ℓ^1 norm) and $\eta = 0$. In the context of compressed sensing, random matrices that are known to fulfill the quotient property with high probability are gaussian, subgaussian, and Weibull matrices [34, 70]. For further details we refer to [12].

3.7 Numerical Results

We conclude this chapter with a series of numerical results. First, in Figures 1 and 2 we show the approximation of several functions via weighted ℓ^1 minimization. Weights of the form $w_i = (u_i)^\alpha$ are used for several different choices of α . In agreement with the discussion in Section 3.2, the choice $\alpha = 1$, i.e., $w_i = u_i$ generally gives among the smallest error. Moreover, while larger values of α sometime give a smaller error, this is not the case for all functions. Notice that in all cases unweighted ℓ^1 minimization gives a worse error than weighted ℓ^1 minimization. As is to be expected, the improvement offered by weighted ℓ^1 minimization in the Chebyshev case is less significant in moderate dimensions than for Legendre polynomials.

The results in Figures 1 and 2 were computed by solving weighted ℓ^1 minimization problems with η set arbitrarily to $\eta = 10^{-12}$ (we make this choice rather than $\eta = 0$ to avoid potential infeasibility issues in the solver). In particular, the condition (33) is not generally satisfied. Following Remark 6, we next assess the size of the constant $Q_u(A)$ defined in (42). Table 1 shows the magnitude of this constant for the setups considered in Figures 1 and 2. Over all ranges of m considered, this constant is never more than 20 in magnitude. That is to say, the additional effect due to the unknown truncation error $\|e_\Lambda\|_2$ is relatively small.

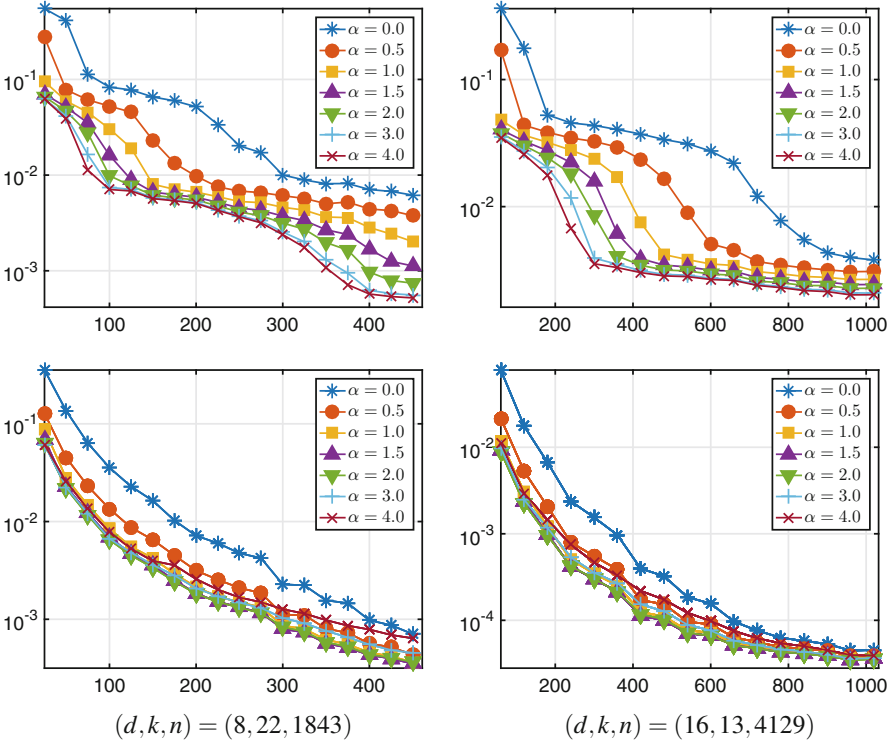


Fig. 1 The error $\|f - \tilde{f}\|_{L^\infty}$ (averaged over 50 trials) against m for Legendre polynomials. Here $\tilde{f} = \sum_{i \in \Lambda} \hat{c}_i \phi_i$, where \hat{c}_Λ is a solution of (11) with weights $w_i = (u_i)^\alpha$ and $\Lambda = \Lambda_k^{\text{HC}}$ a hyperbolic cross index set. The functions used were $f(\mathbf{y}) = \prod_{k=d/2+1}^d \cos(16y_k/2^k) / \prod_{k=1}^{d/2} (1 - y_k/4^k)$ and $f(\mathbf{y}) = \exp\left(-\sum_{k=1}^d y_k/(2d)\right)$ (top and bottom, respectively). The weighted ℓ^1 minimization problem was solved using the SPGL1 package [67, 68] with a maximum of 100,000 iterations and $\eta = 10^{-12}$.

In view of Remark 7, in Figure 3 we assess the performance of weighted ℓ^1 minimization in the presence of external sources of error corrupting the measurements. In order to model this scenario, we consider the problem (11) where the vector of measurements is corrupted by additive noise

$$\mathbf{y} = \frac{1}{\sqrt{m}} (f(z_j))_{j=1}^m + \mathbf{n}, \quad (43)$$

or, equivalently, by recalling (6),

$$\mathbf{y} = \mathbf{A} \mathbf{c}_\Lambda + \mathbf{e}_\Lambda + \mathbf{n}. \quad (44)$$




Fig. 2 The same as Figure 1 but with Chebyshev polynomials.

Table 1 The constant $Q_u(A)$ (averaged over 50 trials) for the setup considered in Figures 1 and 2.

(d, k, n)	m	125	250	375	500	625	750	875	1000
	$(8, 22, 1843)$	Chebyshev	2.65	3.07	3.53	3.95	4.46	5.03	5.78
Legendre		6.45	7.97	8.99	10.5	12.1	13.7	15.8	18.6
$(16, 13, 4129)$	m	250	500	750	1000	1250	1500	1750	2000
	Chebyshev	2.64	2.93	3.30	3.63	3.99	4.41	4.95	5.62
	Legendre	5.64	6.20	6.85	7.60	8.32	8.99	10.1	11.1

We randomly generate the noise as $\mathbf{n} = 10^{-3} \mathbf{g} / \|\mathbf{g}\|_2$, where $\mathbf{g} \in \mathbb{R}^m$ is a standard random gaussian vector, so that $\|\mathbf{n}\|_2 = 10^{-3}$. Considering weights $\mathbf{w} = (u_i^\alpha)_{i \in \Lambda}$, with $\alpha = 0, 1$, we compare the error obtained when the parameter η in (11) is chosen according to each of the following three strategies:

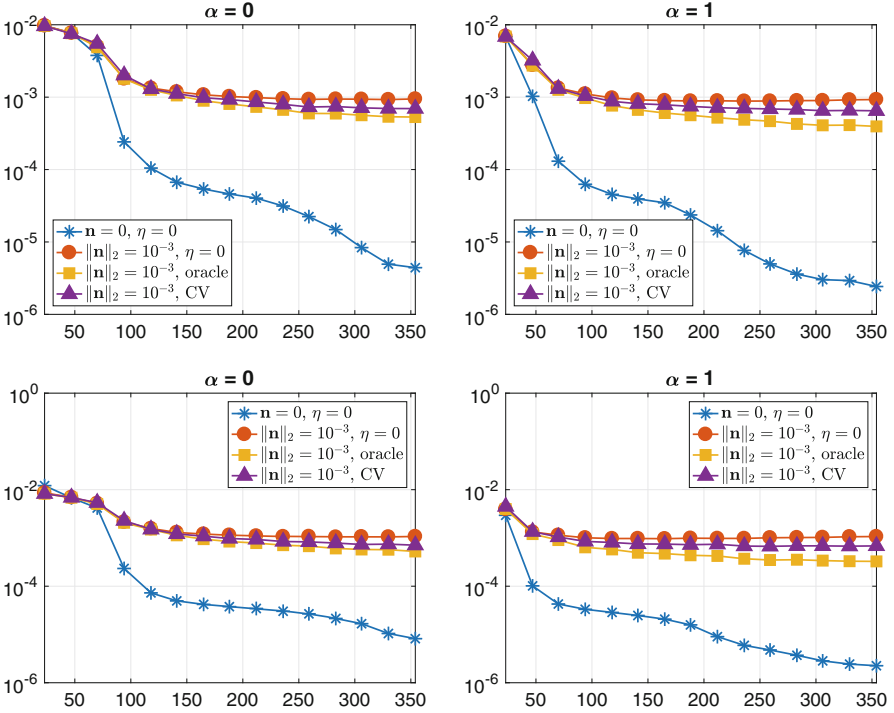


Fig. 3 The error $\|f - \tilde{f}\|_{L^2_v}$ against m . Here $\tilde{f} = \sum_{i \in \Lambda} \hat{c}_i \phi_i$, where \hat{c}_Λ is a solution of (11) with weights $\mathbf{w} = (w_i^c)_{i \in \Lambda}$, with $\alpha = 0$ (left) and $\alpha = 1$ (right), and \mathbf{y} defined as in (43). Regarding $\{\phi_i\}_{i \in \Lambda}$ and ν , the Chebyshev polynomials with the Chebyshev measure are employed in the top line and the Legendre polynomials with the uniform measure in the bottom line. We choose $d = 8$ and $\Lambda = \Lambda_{19}^{\text{HC}}$ with $n = |\Lambda| = 1771$. For each value of m , we average the error over 50 trials considering three different strategies for the choice of η : namely, $\eta = 0$, estimation via oracle least squares, and cross validation (CV). The function approximated is $f(\mathbf{y}) = \exp\left(-\sum_{k=1}^d \cos(y_k)/(8d)\right)$.

1. $\eta = 0$, corresponding to enforcing the exact constraint $\mathbf{A}\mathbf{d} = \mathbf{y}$ in (11);
2. $\eta = \eta_{\text{oracle}} = \|\mathbf{A}\hat{\mathbf{c}}_{\text{oracle}} - \mathbf{y}\|_2$, where $f_{\text{oracle}} = \sum_{i \in \Lambda} (\hat{c}_{\text{oracle}})_i \phi_i$ is the oracle least-squares solution based on $10n$ random samples of f distributed according to ν ;
3. η is estimate using a cross validation approach, as described in [32, Section 3.5], where the search of η is restricted to the values of the form $10^k \cdot \eta_{\text{oracle}}$, where k belongs to a uniform grid of 11 equispaced points on the interval $[-3, 3]$, $3/4$ of the samples are used as reconstruction samples and $1/4$ as validation samples.

The results are in accordance with the estimate (36). Indeed, as expected, for any value of α , the recovery error associated with $\mathbf{n} = \mathbf{0}$ and $\eta = 0$ is always lower than the recovery error associated with $\mathbf{n} \neq \mathbf{0}$ and any choice of η . This can be

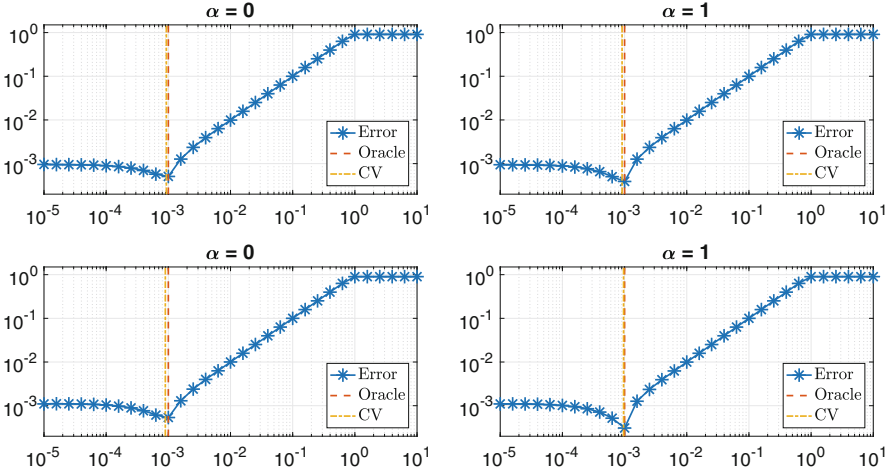


Fig. 4 Recovery error $\|f - \tilde{f}\|_{L^2_\nu}$ (averaged over 50 trials) against η , in the same setting as in Figure 3. We use Chebyshev and Legendre polynomials in the top and bottom rows, respectively. We consider $\eta = 10^k$, with k belonging to a uniform grid of 31 points on the interval $[-5, 1]$. The vertical lines represent the estimated values of η (averaged over 50 trials) based on oracle least squares (red-dashed line) and cross validation (yellow dashed-dotted line). The weights are chosen as $\mathbf{w} = (\mathbf{u}_i^\alpha)_{i \in \Lambda}$, with $\alpha = 0$ (left) and $\alpha = 1$ (right).

explained by the fact that, in the right-hand side of (36), the terms $\sigma_{k,L}(\mathbf{c})/k^{\gamma/2}$ and $\|e_\Lambda\|_2$ are dominated by $\eta + T_u(A, \Lambda, e_\Lambda, \eta)$ when $\mathbf{n} \neq \mathbf{0}$. Moreover, estimating η via oracle least squares (strategy 2) gives better results than cross validation (strategy 3), which in turn is better than the neutral choice $\eta = 0$ (strategy 1). Finally, we note that the discrepancy among the three strategies is accentuated as α gets larger.

In the next experiment we highlight the importance of the parameter η when solving (11) with measurements subject to external sources of error (recall Remark 7). We corrupt the measurements by adding random noise \mathbf{n} with norm $\|\mathbf{n}\|_2 = 10^{-3}$, analogously to (43). Then, for different values of η from 10^{-5} to 10, we solve (11) with weights $\mathbf{w} = (\mathbf{u}_i^\alpha)_{i \in \Lambda}$ and $\alpha = 0, 1$. The resulting recovery errors with respect to the L^2_ν norm (averaged over 50 trials) are plotted as a function of η in Figure 4. For every value of α , the resulting curve is constant for the smallest and largest values of η . In between, the curve exhibits a global minimum, which corresponds to an optimal calibration of η . The values of η estimated via oracle least squares and cross validation are both able to approximate the global minimum on average. However, cross validation has a larger standard deviation compared to the former (see Table 2). This explains why the performance of cross validation is suboptimal in Figure 3. We also notice that the global minimum is more pronounced as α gets larger, in accordance to the observations in Figure 3.

Table 2 Mean \pm standard deviation for the values of η estimated via oracle least squares and cross validation over 50 trials in Figure 4.

α	Chebyshev		Legendre	
	Oracle	Cross validation	Oracle	Cross validation
0	$1.0\text{e-}03 \pm 7.2\text{e-}09$	$9.3\text{e-}04 \pm 3.8\text{e-}04$	$1.0\text{e-}03 \pm 3.7\text{e-}09$	$9.0\text{e-}04 \pm 4.0\text{e-}04$
1	$1.0\text{e-}03 \pm 4.9\text{e-}09$	$9.1\text{e-}04 \pm 4.0\text{e-}04$	$1.0\text{e-}03 \pm 3.6\text{e-}09$	$9.7\text{e-}04 \pm 3.6\text{e-}04$

4 Conclusions and Challenges

The concern of this chapter has been the emerging topic of compressed sensing for high-dimensional approximation. As shown, smooth, multivariate functions are compressible in orthogonal polynomial bases. Moreover, their coefficients have a certain form of structured sparsity corresponding to so-called lower sets. The main result of this work is that such structure can be exploited via weighted ℓ^1 -norm regularizers. Doing so leads to sample complexity estimates that are at most logarithmically dependent on the dimension d , thus mitigating the curse of dimensionality to a substantial extent.

As discussed in Section 1.5, this topic has garnered much interest over the last half a dozen years. Yet challenges remain. We conclude by highlighting a number of open problems in this area:

Unbounded domains We have considered only bounded hypercubes in this chapter. The case of unbounded domains presents additional issues. While Hermite polynomials (orthogonal on \mathbb{R}) have been considered in the case of unweighted ℓ^1 minimization in [39, 41, 53], the corresponding measurement conditions exhibit exponentially large factors in either the dimension d or degree k of the (total degree) index space used. It is unclear how to obtain dimension-independent measurement conditions in this setting, even for structured sparsity in lower sets.

Sampling strategies Throughout we have considered sampling i.i.d. according to the orthogonality measure of the basis functions. This is by no means the only choice, and various other sampling strategies have been considered in other works [39, 41, 44, 52, 53, 64, 71]. Empirically, several of these approaches are known to give some benefits. However, it is not known how to design sampling strategies which lead to better measurement conditions than those given in Theorem 5. A singular challenge is to design a sampling strategy for which m need only scale linearly with k . We note in passing that this has been achieved for the oracle least-squares estimator (recall Section 3.3) [28]. However, it is not clear how to extend this approach to a compressed sensing framework.

Alternatives to weighted ℓ^1 minimization. As discussed in Remark 1, lower set structure is a type of structured sparsity model. We have used weighted ℓ^1 minimization to promote such structure. Yet other approaches may convey benefits. Different, but related, types of structured sparsity have been exploited in the past

using greedy or iterative algorithms [5, 9, 30, 33], or by designing appropriate convex regularizers [66]. This remains an interesting problem for future work.

Recovering Hilbert-valued functions. We have focused on compressed sensing-based polynomial approximation of high-dimensional functions whose coefficients belong to the complex domain \mathbb{C} . However, an important problem in computational science, especially in the context of UQ and optimal control, involves the approximation of parametric PDEs. Current compressed sensing techniques proposed in literature [10, 22, 32, 47, 56, 59, 73] only approximate functionals of parameterized solutions, e.g., evaluation at a single spatial location, whereas a more robust approach should consider an ℓ_1 -regularized problem involving Hilbert-valued signals, i.e., signals where each coordinate is a function in a Hilbert space, which can provide a direct, global reconstruction of the solutions in the entire physical domain. However, to achieve this goal new iterative minimization procedures as well as several theoretical concepts will need to be extended to the Hilbert space setting. The advantages of this approach over pointwise recovery with standard techniques will include: (i) for many parametric and stochastic model problems, global estimate of solutions in the physical domain is a quantity of interest; (ii) the recovery guarantees of this strategy can be derived from the decay of the polynomial coefficients in the relevant function space, which is well-known in the existing theory; and (iii) the global reconstruction only assumes *a priori* bounds of the tail expansion in energy norms, which are much more realistic than pointwise bounds.

Acknowledgements The first and second authors acknowledge the support of the Alfred P. Sloan Foundation and the Natural Sciences and Engineering Research Council of Canada through grant 611675. The second author acknowledges the Postdoctoral Training Center in Stochastics of the Pacific Institute for the Mathematical Sciences for the support. The third author acknowledges support by the US Defense Advanced Research Projects Agency, Defense Sciences Office under contract and award numbers HR0011619523 and 1868-A017-15; the US Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under contract number ERKJ259 and ERKJ314; and the Laboratory Directed Research and Development program at the Oak Ridge National Laboratory, which is operated by UT-Battelle, LLC., for the US Department of Energy under Contract DE-AC05-00OR22725.

References

1. B. Adcock, Infinite-dimensional compressed sensing and function interpolation. *Found. Comput. Math.*, 1–41 (2017). <https://doi.org/10.1007/s10208-017-9350-3>
2. B. Adcock, Infinite-dimensional ℓ^1 minimization and function approximation from pointwise data. *Constr. Approx.* **45**(3), 345–390 (2017)
3. B. Adcock, A. Bao, S. Brugiapaglia, Correcting for unknown errors in sparse high-dimensional function approximation (2017). arXiv:1711.07622
4. B. Adcock, A.C. Hansen. Generalized sampling and infinite-dimensional compressed sensing. *Found. Comput. Math.* **16**(5), 1263–1323 (2016)

5. R.G. Baraniuk, V. Cevher, M.F. Duarte, C. Hedge, Model-based compressive sensing. *IEEE Trans. Inform. Theory* **56**(4), 1982–2001 (2010)
6. J. Beck, F. Nobile, L. Tamellini, R. Tempone, Convergence of quasi-optimal Stochastic Galerkin methods for a class of PDEs with random coefficients. *Comput. Math. Appl.* **67**(4), 732–751 (2014)
7. R.E. Bellman, *Adaptive Control Processes: A Guided Tour* (Princeton University Press, Princeton, 1961)
8. J. Bigot, C. Boyer, P. Weiss, An analysis of block sampling strategies in compressed sensing. *IEEE Trans. Inform. Theory* **64**(4), 2125–2139 (2016)
9. T. Blumensath, Sampling theorems for signals from the union of finite-dimensional linear subspaces. *IEEE Trans. Inform. Theory* **55**(4), 1872–1882 (2009)
10. J.-L. Bouchot, H. Rauhut, C. Schwab, Multi-level Compressed Sensing Petrov-Galerkin discretization of high-dimensional parametric PDEs (2017). arXiv:1701.01671
11. S. Brugiapaglia, COMpRessed SolvING: sparse approximation of PDEs based on compressed sensing, Ph.D. thesis, Politecnico di Milano, Milano, 2016
12. S. Brugiapaglia, B. Adcock, Robustness to unknown error in sparse regularization (2017). arXiv:1705.10299
13. S. Brugiapaglia, F. Nobile, S. Micheletti, S. Perotto, A theoretical study of compressed solving for advection-diffusion-reaction problems. *Math. Comput.* **87**(309), 1–38 (2018)
14. H.-J. Bungartz, M. Griebel, Sparse grids. *Acta Numer.* **13**, 147–269 (2004)
15. E.J. Candès, Y. Plan, A probabilistic and RIPless theory of compressed sensing. *IEEE Trans. Inform. Theory* **57**(11), 7235–7254 (2011)
16. E.J. Candès, J. Romberg, T. Tao, Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory* **52**(1), 489–509 (2006)
17. A. Chernov, D. Düng, New explicit-in-dimension estimates for the cardinality of high-dimensional hyperbolic crosses and approximation of functions having mixed smoothness. *J. Compl.* **32**, 92–121 (2016)
18. A. Chkifa, A. Cohen, R. DeVore, C. Schwab, Sparse adaptive Taylor approximation algorithms for parametric and stochastic elliptic PDEs. *Modél. Math. Anal. Numér.* **47**(1), 253–280 (2013)
19. A. Chkifa, A. Cohen, G. Migliorati, F. Nobile, R. Tempone, Discrete least squares polynomial approximation with random evaluations – application to parametric and stochastic elliptic PDEs. *ESAIM Math. Model. Numer. Anal.* **49**(3), 815–837 (2015)
20. A. Chkifa, A. Cohen, C. Schwab, High-dimensional adaptive sparse polynomial interpolation and applications to parametric PDEs. *Found. Comput. Math.* **14**(4), 601–633 (2014)
21. A. Chkifa, A. Cohen, C. Schwab, Breaking the curse of dimensionality in sparse polynomial approximation of parametric PDEs. *J. Math. Pures Appl.* **103**, 400–428 (2015)
22. A. Chkifa, N. Dexter, H. Tran, C.G. Webster, Polynomial approximation via compressed sensing of high-dimensional functions on lower sets. *Math. Comput.* arXiv:1602.05823 (2016, to appear)
23. I.-Y. Chun, B. Adcock, Compressed sensing and parallel acquisition. *IEEE Trans. Inform. Theory* **63**(8), 4760–4882 (2017). arXiv:1601.06214
24. A. Cohen, M.A. Davenport, D. Leviatan, On the stability and accuracy of least squares approximations. *Found. Comput. Math.* **13**, 819–834 (2013)
25. A. Cohen, R. DeVore, Approximation of high-dimensional parametric PDEs. *Acta Numer.* **24**, 1–159 (2015)
26. A. Cohen, R.A. DeVore, C. Schwab, Convergence rates of best N -term Galerkin approximations for a class of elliptic sPDEs. *Found. Comput. Math.* **10**, 615–646 (2010)
27. A. Cohen, R.A. DeVore, C. Schwab, Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDEs. *Anal. Appl.* **9**(1), 11–47 (2011)
28. A. Cohen, G. Migliorati, Optimal weighted least-squares methods (2016). arXiv:1608.00512
29. A. Cohen, G. Migliorati, F. Nobile, Discrete least-squares approximations over optimized downward closed polynomial spaces in arbitrary dimension. *Constr. Approx.* **45**(3), 497–519 (2017)

30. M.A. Davenport, M.F. Duarte, Y.C. Eldar, G. Kutyniok, Introduction to compressed sensing, in *Compressed Sensing: Theory and Applications* (Cambridge University Press, Cambridge, 2011)
31. D.L. Donoho, Compressed sensing. *IEEE Trans. Inform. Theory* **52**(4), 1289–1306 (2006)
32. A. Doostan, H. Owhadi, A non-adapted sparse approximation of PDEs with stochastic inputs. *J. Comput. Phys.* **230**(8), 3015–3034 (2011)
33. M.F. Duarte, Y.C. Eldar, Structured compressed sensing: from theory to applications. *IEEE Trans. Signal Process.* **59**(9), 4053–4085 (2011)
34. S. Foucart, Stability and robustness of ℓ_1 -minimizations with weibull matrices and redundant dictionaries. *Linear Algebra Appl.* **441**, 4–21 (2014)
35. S. Foucart, H. Rauhut, *A Mathematical Introduction to Compressive Sensing* (Birkhauser, Basel, 2013)
36. D. Gross, Recovering low-rank matrices from few coefficients in any basis. *IEEE Trans. Inform. Theory* **57**(3), 1548–1566 (2011)
37. M. Gunzburger, C.G. Webster, G. Zhang, Stochastic finite element methods for partial differential equations with random input data. *Acta Numer.* **23**, 521–650 (2014)
38. M. Gunzburger, C.G. Webster, G. Zhang, Sparse collocation methods for stochastic interpolation and quadrature, in *Handbook of Uncertainty Quantification* (Springer, New York, 2016), pp. 1–46
39. L. Guo, A. Narayan, T. Zhou, Y. Chen, Stochastic collocation methods via L1 minimization using randomized quadratures. *SIAM J. Sci. Comput.* **39**(1), A333–A359 (2017). arXiv:1602.00995
40. J. Hampton, A. Doostan, Coherence motivated sampling and convergence analysis of least squares polynomial Chaos regression. *Comput. Methods Appl. Mech. Eng.* **290**, 73–97 (2015)
41. J. Hampton, A. Doostan, Compressive sampling of polynomial chaos expansions: convergence analysis and sampling strategies. *J. Comput. Phys.* **280**, 363–386 (2015)
42. V.H. Hoang, C. Schwab, Regularity and generalized polynomial chaos approximation of parametric and random 2nd order hyperbolic partial differential equations. *Anal. Appl.* **10**(3), 295–326 (2012)
43. J.D. Jakeman, M.S. Eldred, K. Sargsyan, Enhancing l_1 -minimization estimates of polynomial chaos expansions using basis selection. *J. Comput. Phys.* **289**, 18–34 (2015). arXiv:1407.8093
44. J.D. Jakeman, A. Narayan, T. Zhou, A generalized sampling and preconditioning scheme for sparse approximation of polynomial chaos expansions. *SIAM J. Sci. Comput.* **39**(3), A1114–A1144 (2017). arXiv:1602.06879
45. T. Kühn, W. Sickel, T. Ullrich, Approximation of mixed order Sobolev functions on the d -torus: asymptotics, preasymptotics, and d -dependence. *Constr. Approx.* **42**(3), 353–398 (2015)
46. O.P. Le Maître, O.M. Knio, *Spectral Methods for Uncertainty Quantification* (Springer, New York, 2010)
47. L. Mathelin, K.A. Gallivan, A compressed sensing approach for partial differential equations with random input data. *Commun. Comput. Phys.* **12**(4), 919–954 (2012)
48. G. Migliorati, Polynomial approximation by means of the random discrete L^2 projection and application to inverse problems for PDEs with stochastic data, Ph.D. thesis, Politecnico di Milano, Milano, 2013
49. G. Migliorati, Multivariate Markov-type and Nikolskii-type inequalities for polynomials associated with downward closed multi-index sets. *J. Approx. Theory* **189**, 137–159 (2015)
50. G. Migliorati, F. Nobile, Analysis of discrete least squares on multivariate polynomial spaces with evaluations at low-discrepancy point sets. *J. Complexity* **31**(4), 517–542 (2015)
51. G. Migliorati, F. Nobile, E. von Schwerin, R. Tempone, Analysis of the discrete L^2 projection on polynomial spaces with random evaluations. *Found. Comput. Math.* **14**, 419–456 (2014)
52. A. Narayan, T. Zhou, Stochastic collocation on unstructured multivariate meshes. *Commun. Comput. Phys.* **18**(1), 1–36 (2015)
53. A. Narayan, J.D. Jakeman, T. Zhou, A Christoffel function weighted least squares algorithm for collocation approximations. *Math. Comput.* **86**(306), 1913–1947 (2014). arXiv:1412.4305

54. F. Nobile, R. Tempone, C.G. Webster, An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.* **46**(5), 2411–2442 (2008)
55. F. Nobile, R. Tempone, C.G. Webster, A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.* **46**(5), 2309–2345 (2008)
56. J. Peng, J. Hampton, A. Doostan, A weighted ℓ_1 -minimization approach for sparse polynomial chaos expansions. *J. Comput. Phys.* **267**, 92–111 (2014)
57. J. Peng, J. Hampton, A. Doostan, On polynomial chaos expansion via gradient-enhanced ℓ_1 -minimization. *J. Comput. Phys.* **310**, 440–458 (2016)
58. H. Rauhut, Random sampling of sparse trigonometric polynomials. *Appl. Comput. Harmon. Anal.* **22**(1), 16–42 (2007)
59. H. Rauhut, C. Schwab, Compressive sensing Petrov-Galerkin approximation of high dimensional parametric operator equations. *Math. Comput.* **86**, 661–700 (2017)
60. H. Rauhut, R. Ward, Sparse Legendre expansions via ℓ_1 -minimization. *J. Approx. Theory* **164**(5), 517–533 (2012)
61. H. Rauhut, R. Ward, Interpolation via weighted ℓ_1 minimization. *Appl. Comput. Harmon. Anal.* **40**(2), 321–351 (2016)
62. M.K. Stoyanov, C.G. Webster, A dynamically adaptive sparse grid method for quasi-optimal interpolation of multidimensional functions. *Comput. Math. Appl.* **71**(11), 2449–2465 (2016)
63. G. Szegő, *Orthogonal Polynomials* (American Mathematical Society, Providence, RI, 1975)
64. G. Tang, G. Iaccarino, Subsampled Gauss quadrature nodes for estimating polynomial chaos expansions. *SIAM/ASA J. Uncertain. Quantif.* **2**(1), 423–443 (2014)
65. H. Tran, C.G. Webster, G. Zhang, Analysis of quasi-optimal polynomial approximations for parameterized PDEs with deterministic and stochastic coefficients. *Numer. Math.* **137**(2), 451–493 (2017). arXiv:1508.01821
66. Y. Traonmilin, R. Gribonval, Stable recovery of low-dimensional cones in Hilbert spaces: one RIP to rule them all. *Appl. Comput. Harm. Anal.* (2017). <https://doi.org/10.1016/j.acha.2016.08.004>
67. E. van den Berg, M.P. Friedlander, SPGL1: a solver for large-scale sparse reconstruction (June 2007), <http://www.cs.ubc.ca/labs/scl/spgl1>
68. E. van den Berg, M.P. Friedlander, Probing the Pareto frontier for basis pursuit solutions. *SIAM J. Sci. Comput.* **31**(2), 890–912 (2008)
69. C.G. Webster, Sparse grid stochastic collocation techniques for the numerical solution of partial differential equations with random input data, Ph.D. thesis, Florida State University, Tallahassee, 2007
70. P. Wojtaszczyk, Stability and instance optimality for gaussian measurements in compressed sensing. *Found. Comput. Math.* **10**(1), 1–13 (2010)
71. Z. Xu, T. Zhou, On sparse interpolation and the design of deterministic interpolation points. *SIAM J. Sci. Comput.* **36**(4), 1752–1769 (2014)
72. L. Yan, L. Guo, D. Xiu, Stochastic collocation algorithms using ℓ_1 -minimization. *Int. J. Uncertain. Quantif.* **2**(3), 279–293 (2012)
73. X. Yang, G.E. Karniadakis, Reweighted ℓ_1 minimization method for stochastic elliptic differential equations. *J. Comput. Phys.* **248**, 87–108 (2013)
74. X. Yang, H. Lei, N.A. Baker, G. Lin, Enhancing sparsity of Hermite polynomial expansions by iterative rotations. *J. Comput. Phys.* **307**, 94–109 (2016). arXiv:1506.04344