

Correcting for unknown errors in sparse high-dimensional function approximation

Ben Adcock*, Anyi Bao† and Simone Brugiapaglia‡

May 10, 2019

Abstract

We consider sparsity-based techniques for the approximation of high-dimensional functions from random pointwise evaluations. To date, almost all the works published in this field contain some *a priori* assumptions about the error corrupting the samples that are hard to verify in practice. In this paper, we instead focus on the scenario where the error is unknown. We study the performance of four sparsity-promoting optimization problems: weighted quadratically-constrained basis pursuit, weighted LASSO, weighted square-root LASSO, and weighted LAD-LASSO. From the theoretical perspective, we prove uniform recovery guarantees for these decoders, deriving recipes for the optimal choice of the respective tuning parameters. On the numerical side, we compare them in the pure function approximation case and in applications to uncertainty quantification of ODEs and PDEs with random inputs. Our main conclusion is that the lesser-known square-root LASSO is better suited for high-dimensional approximation than the other procedures in the case of bounded noise, since it avoids (both theoretically and numerically) the need for parameter tuning.

1 Introduction

Sparse regularization has been recently proved to be a useful tool for the approximation of functions defined over d -dimensional domains, due to its ability to lessen the curse of dimensionality when combined with compressed sensing principles. This approach is based on three main elements: sparse approximation of the function with respect to orthogonal polynomials, random pointwise sampling, and sparse recovery via a (weighted) ℓ^1 minimization decoder. Combining these three ingredients, it is possible to construct approximations using a number of pointwise evaluations that depends only logarithmically on d [1, 4, 23]. This feature is particularly appealing in uncertainty quantification applications, where the function to approximate is a quantity of interest of a parametric PDE [48, 67].

This paper concerns the effective treatment of errors in the data by sparse recovery procedures. There are three primary sources of such error:

- (i) *Truncation error.* This occurs when infinite expansion in an orthonormal polynomial basis is replaced by a finite expansion. This is an intrinsic source of error in high-dimensional function approximation, which does not depend on the particular application considered.
- (ii) *Discretization error.* In practice, in applications to parametric PDEs, function samples are computed using a numerical PDE solver. This error arises due to the discretization in the numerical routine (e.g. introducing a finite element mesh or performing a Galerkin projection).
- (iii) *Numerical error.* Following (ii), this further source of error occurs when the discretized problem is solved numerically (e.g. via an iterative solver).

*Simon Fraser University, Burnaby, BC, Canada. e-mail: ben_adcock@sfu.ca

†University of British Columbia, Vancouver, BC, Canada. e-mail: abao@math.ubc.ca

‡Simon Fraser University, Burnaby, BC, Canada. e-mail: simone_brugiapaglia@sfu.ca

To date, almost all works in sparse regularization for high-dimensional approximation have assumed an *a priori* bound for this error. Specifically, if the vector of data is $\mathbf{y} = \mathbf{y}^{\text{exact}} + \mathbf{e}$, where $\mathbf{y}^{\text{exact}}$ is the error-free data and \mathbf{e} is the vector of errors, then one assumes the bound

$$\|\mathbf{e}\|_2 \leq \eta, \quad (1)$$

for some known $\eta > 0$. In this case, sparse regularization performed using the (weighted) quadratically-constrained basis pursuit decoder admits rigorous theoretical recovery guarantees [1, 4, 23, 51, 66]. In practice, however, a bound of the form (1) is usually *unknown*, since the sources of error (i), (ii), and (iii) are function dependent. Cross validation is a common empirical remedy to tune such a parameter [29, 40, 67]. Yet this is computationally expensive, and lacks theoretical guarantees. We also observe that for certain model problems (in the context of parametric PDEs) theoretical estimates for η are available. However, the presence of implicit constants can create numerical issues, as we will demonstrate in our numerical experiments.

In this paper, we study three different decoders for sparse recovery: weighted quadratically-constrained basis pursuit, weighted LASSO and weighted square-root LASSO. In all cases, we consider their performance for *unknown errors*: that is, in the absence of the bound (1). These decoders are compared from both the theoretical and numerical viewpoints. On the theoretical side, we prove uniform and robust recovery error estimates for the proposed decoders. From the numerical viewpoint, we compare their recovery performances in the case of pure function approximation and in the case of parametric ODEs and PDEs.

One of our main messages is that the lesser-known weighted square-root LASSO is more suitable for sparse high-dimensional function approximation than the commonly-used quadratically-constrained basis pursuit decoder. It exhibits recovery performances comparable (if not superior) to the others, but the optimal choice of its tuning parameter does not depend on the noise level in the data. Therefore, it does not require any *a priori* knowledge on the error corrupting the samples. This fact is well-known by statisticians; however, it does not seem to have been fully exploited by the high-dimensional function approximation and the compressed sensing communities.

To complete the paper, we also consider a fourth type of error. In addition to (i), (ii), and (iii), in large-scale uncertainty quantification implementations it is increasingly common to encounter

- (iv) *Corruption error*. This occurs when an unknown subset of the measurements are corrupted arbitrarily, for instance due to a node failure when computations are performed on a cluster.

This is a fundamentally different type of error to the others. Corruption errors can be arbitrarily large, but only affect a (typically unknown) fraction of the measurements. They arise in modern complex computational frameworks consisting of many interconnected parallel components, which are subject to faults. These faults may be hard to detect and heavily pollute the measurements [14].

In this setting, the vector of errors is of the form $\mathbf{e} = \mathbf{e}^{\text{bounded}} + \mathbf{e}^{\text{sparse}}$, where $\mathbf{e}^{\text{bounded}}$ is a combination of (i), (ii), or (iii) and has small norm, and $\mathbf{e}^{\text{sparse}}$ is of type (iv) and it is sparse with large norm. Methods that are robust to this type of errors are often called *fault-tolerant* or *resilient* and have already been considered in uncertainty quantification and in compressed sensing [3, 41, 52, 63]. To correct for this type of error, we study a fourth decoder; the weighted LAD-LASSO. Numerically, we show that this method can achieve quite striking performance; in our experiments, it can easily correct for over 10% corrupted measurements. Much like the square-root LASSO, we conclude that this should be the method of choice in the presence of errors (i), (ii), (iii), and (iv).

The paper is organized as follows. In §2, we describe the problem setting, defining the high-dimensional function approximation framework and the four decoders considered for weighted ℓ^1 minimization; moreover, we present the main contributions of the paper, i.e. uniform recovery guarantees for the proposed decoders under unknown error corrupting the measurements. §3 contains a review of the literature for each decoder. In §4, we compare the four decoders from a numerical perspective in the case of pure function approximation and in applications to parametric PDEs and ODEs with random inputs. §5 deals with the theoretical analysis; this is the mathematical core of the paper and it is the most technical part. Finally, in §6 we give our conclusions and state some open problems.

2 Main contributions

We now present the main contributions of the paper. We define the sparse high-dimensional approximation methodology in §2.1. Then, in §2.3 we introduce the concept of sparsity in lower sets and other related notions, which are necessary to understand the theoretical analysis. Finally, in §2.4 we state the uniform robust recovery error estimates for the four decoders (proofs of these results are given in §5).

2.1 Sparse high-dimensional approximation

Let $d \in \mathbb{N} := \{1, 2, 3, \dots\}$, with $d \gg 1$, and consider a high-dimensional complex-valued function

$$f : D \rightarrow \mathbb{C}, \quad \text{with } D := (-1, 1)^d.$$

We define d probability measures $\nu^{(1)}, \dots, \nu^{(d)}$ on $(-1, 1)$ and the corresponding tensor product measure on D as

$$\nu := \nu^{(1)} \otimes \dots \otimes \nu^{(d)}.$$

Moreover, we consider d orthonormal bases $\{\phi_j^{(\ell)}\}_{j \in \mathbb{N}_0}$ of $L^2_{\nu^{(\ell)}}(-1, 1)$ for $\ell \in [d]$, where $[d] := \{1, \dots, d\}$ and $\mathbb{N}_0 := \{0\} \cup \mathbb{N}$. Then, we build $\{\phi_{\mathbf{i}}\}_{\mathbf{i} \in \mathbb{N}_0^d}$ as the corresponding tensor product basis of $L^2_\nu(D)$:

$$\phi_{\mathbf{i}} := \phi_{i_1}^{(1)} \otimes \dots \otimes \phi_{i_d}^{(d)}, \quad \forall \mathbf{i} \in \mathbb{N}_0^d. \quad (2)$$

Assume $f \in L^2_\nu(D) \cap L^\infty(D)$ and consider its expansion

$$f = \sum_{\mathbf{i} \in \mathbb{N}_0^d} x_{\mathbf{i}} \phi_{\mathbf{i}}, \quad (3)$$

where $x_{\mathbf{i}} = \langle f, \phi_{\mathbf{i}} \rangle_{L^2_\nu}$, for every $\mathbf{i} \in \mathbb{N}_0^d$. Throughout the paper, we define the sequence of coefficients

$$\mathbf{x} := (x_{\mathbf{i}})_{\mathbf{i} \in \mathbb{N}_0^d} \in \ell^2(\mathbb{N}_0^d).$$

where $\ell^p(\mathbb{N}_0^d)$ denotes the space of sequences with finite ℓ^p norm. Given a subset $\Lambda \subseteq \mathbb{N}_0^d$ of finite cardinality $|\Lambda| = n$, we obtain a finite approximation of f by considering the truncated basis $\{\phi_{\mathbf{i}}\}_{\mathbf{i} \in \Lambda}$. The corresponding finite-dimensional vector of coefficients is denoted by \mathbf{x}_Λ . Assuming the set of multi-indices in Λ is ordered as $\mathbf{i}_1, \dots, \mathbf{i}_n$, we can also write

$$f_\Lambda := \sum_{j=1}^n x_{\mathbf{i}_j} \phi_{\mathbf{i}_j} \quad \text{and} \quad \mathbf{x}_\Lambda := (x_{\mathbf{i}_j})_{j \in [n]}, \quad (4)$$

The vector \mathbf{x}_Λ will be equivalently considered as an element of \mathbb{C}^n or of $\ell^2(\mathbb{N}_0^d)$. In the latter case, the components $x_{\mathbf{i}}$ with $\mathbf{i} \notin \Lambda$ are implicitly assumed to be zero.

We aim at recovering a sparse approximation of f with respect to the basis $\{\phi_{\mathbf{i}}\}_{\mathbf{i} \in \Lambda}$ from pointwise samples. Hence, we define the *vector of samples* $\mathbf{y} \in \mathbb{C}^m$ and the *design matrix* $\mathbf{A} \in \mathbb{C}^{m \times n}$ as¹

$$y_i := \frac{1}{\sqrt{m}}(f(\mathbf{t}_i)), \quad A_{ij} := \frac{1}{\sqrt{m}}(\phi_{\mathbf{i}_j}(\mathbf{t}_i)), \quad \forall i \in [m], j \in [n], \quad (5)$$

where, as in previous works in high-dimensional polynomial approximation [23, 29, 67], the samples are distributed independently at random according to the orthogonality measure ν , i.e.

$$\mathbf{t}_1, \dots, \mathbf{t}_m \stackrel{\text{i.i.d.}}{\sim} \nu(\mathbf{t}). \quad (6)$$

Then, the truncated vector of coefficients $\mathbf{x}_\Lambda \in \mathbb{C}^n$ satisfies

$$\mathbf{y} = \mathbf{A} \mathbf{x}_\Lambda + \mathbf{e}, \quad (7)$$

¹The factor $1/\sqrt{m}$ is needed in order to guarantee the restricted isometry property for the design matrix \mathbf{A} . See §5.

where $\mathbf{e} \in \mathbb{C}^m$ is an unknown source of error corrupting the samples. The vector \mathbf{e} contains (at the very least) the truncation error

$$(e_{\Lambda}^{\text{trunc}})_i := \frac{1}{\sqrt{m}}(f(\mathbf{t}_i) - f_{\Lambda}(\mathbf{t}_i)) = \frac{1}{\sqrt{m}} \sum_{\mathbf{i} \notin \Lambda} x_{\mathbf{i}} \phi_{\mathbf{i}}(\mathbf{t}_i), \quad \forall i \in [m]. \quad (8)$$

In general, $\mathbf{e} = \mathbf{e}_{\Lambda}^{\text{trunc}} + \mathbf{e}^{\text{misc}}$, where \mathbf{e}^{misc} contains other sources of error, such as discretization error, numerical error, faults, or other forms of noise, all of which occur commonly in uncertainty quantification applications. In this paper we shall consider the following two types of noise models:

- (a) *Bounded noise.* Here $\|\mathbf{e}\|_2$ is assumed to be small and the goal is to recover f up to this error. This is the case for truncation and numerical errors (types (i), (ii), and (iii) in §1), for instance.
- (b) *Bounded + unbounded sparse noise.* In this case, $\mathbf{e} = \mathbf{e}^{\text{bounded}} + \mathbf{e}^{\text{sparse}}$, where $\mathbf{e}^{\text{bounded}}$ is as in (a) and $\mathbf{e}^{\text{sparse}}$ has only a small number of nonzero entries, but which can be arbitrarily large and whose locations are unknown. This is the case when corruption errors, due for instance to solver faults, are additionally present (type (iv) in §1). The goal in this case is to correct for the corruption error, and to recover f up to $\|\mathbf{e}^{\text{bounded}}\|_2$.

Note that, since the function to be approximated is high-dimensional, the cardinality n of the truncated index set Λ is typically large. Therefore, we focus on the regime $m < n$, where (7) is underdetermined.

Finally, we will consider two particular examples of orthonormal systems of $L_{\nu}^2(D)$: the well-known tensorized Legendre and Chebyshev polynomials. In d dimensions, these are orthogonal with respect to the tensor uniform and Chebyshev measures

$$d\nu = \frac{1}{2^d} d\mathbf{t} \quad (\text{Legendre}), \quad d\nu = \prod_{\ell=1}^d \frac{1}{\pi \sqrt{1-t_{\ell}^2}} d\mathbf{t} \quad (\text{Chebyshev}).$$

In each case, the samples $\mathbf{t}_1, \dots, \mathbf{t}_m$ are randomly and independently distributed according to ν .

2.2 Decoders

In order to recover an approximate solution $\hat{\mathbf{x}}_{\Lambda} = (\hat{x}_{\mathbf{i}_j})_{j \in [n]} \in \mathbb{C}^n$ from (7), we use weighted ℓ^1 minimization. Given a sequence of *positive* weights $\mathbf{w} = (w_{\mathbf{i}})_{\mathbf{i} \in \mathbb{N}_0^d}$, the corresponding weighted ℓ^p norm, referred to as the $\ell_{\mathbf{w}}^p$ norm, is defined as

$$\|\mathbf{x}\|_{p, \mathbf{w}} := \left(\sum_{\mathbf{i} \in \mathbb{N}_0^d} w_{\mathbf{i}}^{2-p} |x_{\mathbf{i}}|^p \right)^{1/p},$$

and $\ell_{\mathbf{w}}^p(\mathbb{N}_0^d)$ is the space of sequences having finite $\ell_{\mathbf{w}}^p$ norm.

Given weights \mathbf{u} and errors of the form (i) and (ii), and (iii) (see §1), we compute $\hat{\mathbf{x}}_{\Lambda}$ with one of the following decoders:

Weighted Quadratically-Constrained Basis Pursuit (WQCBP)

$$\hat{\mathbf{x}}_{\Lambda} = \arg \min_{\mathbf{z} \in \mathbb{C}^n} \|\mathbf{z}\|_{1, \mathbf{u}} \text{ s.t. } \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2 \leq \eta, \quad (9)$$

Weighted LASSO (WLASSO)

$$\hat{\mathbf{x}}_{\Lambda} = \arg \min_{\mathbf{z} \in \mathbb{C}^n} \|\mathbf{z}\|_{1, \mathbf{u}} + \lambda \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2, \quad (10)$$

Weighted Square-Root LASSO (WSR-LASSO)

$$\hat{\mathbf{x}}_{\Lambda} = \arg \min_{\mathbf{z} \in \mathbb{C}^n} \|\mathbf{z}\|_{1, \mathbf{u}} + \lambda \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2, \quad (11)$$

The above decoders are designed to deal specifically with bounded noise (model (a) in §2.1). Conversely, if the measurements additionally contain corruption errors (model (b) in §2.1), in which case \mathbf{e} has a small number of large entries, then we instead consider the following decoder:

Weighted LAD-LASSO (WLAD-LASSO)

$$\hat{\mathbf{x}}_\Lambda = \arg \min_{\mathbf{z} \in \mathbb{C}^n} \|\mathbf{z}\|_{1,\mathbf{u}} + \lambda \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_1. \quad (12)$$

The use of the ℓ^1 -norm in the fitting term is to exploit the sparsity of the corruptions in this case.

Notice that the minimizer $\hat{\mathbf{x}}_\Lambda$ need not be necessarily unique in our framework. Given $\hat{\mathbf{x}}_\Lambda$, the resulting approximation \tilde{f} of f is then defined by

$$\tilde{f} := \sum_{\mathbf{i} \in \Lambda} \hat{x}_{\mathbf{i}} \phi_{\mathbf{i}}. \quad (13)$$

We assume throughout that $\eta \geq 0$ and $\lambda > 0$. These will be referred to as *tuning parameters*. When $\eta = 0$, the optimization program (9) is also referred to as *weighted basis pursuit (WBP)*. It is worth noting that WQCBP is the only optimization program in constrained form. We will comment more on these decoders while reviewing the literature (§3).

As for the choice of the weights, \mathbf{u} is chosen as follows:

$$u_{\mathbf{i}} = \|\phi_{\mathbf{i}}\|_{L^\infty}, \quad \forall \mathbf{i} \in \mathbb{N}_0^d. \quad (14)$$

They are referred to as *intrinsic weights* [1]. Notice that $u_{\mathbf{i}} := \|\phi_{\mathbf{i}}\|_{L^\infty} \geq \|\phi_{\mathbf{i}}\|_{L_v^2} = 1$. This particular choice of weights has been proved to be effective for WQCBP, both theoretically and numerically [1, 4, 23]. For the polynomial bases employed here, we have

$$u_{\mathbf{i}} = \prod_{\ell=1}^d \sqrt{2i_\ell + 1} \quad (\text{Legendre}), \quad u_{\mathbf{i}} = 2^{\|\mathbf{i}\|_0/2} \quad (\text{Chebyshev}). \quad (15)$$

2.3 Sparsity in lower sets

Compressed sensing theory considers the recovery of sparse vectors from limited numbers of measurements. Many works have sought to apply these principles to high-dimensional approximation (see §3). Unfortunately, standard application of these principles leads to sample complexities that depend exponentially on the dimension d (see, for example, [4]). More recently [1, 23], it has been shown that this *curse of dimensionality in the sample complexity* can be overcome by considering certain structured sparsity models based on so-called *lower sets* (also known as *downward closed* or *monotone sets*):

Definition 2.1 (Lower set). The set $S \subseteq \mathbb{N}_0^d$ is said to be a *lower set* if whenever $\mathbf{i} \in S$ and $\mathbf{j} \leq \mathbf{i}$ (where the inequality holds componentwise) then $\mathbf{j} \in S$.

Employing this definition, we consider the space of s -sparse vectors whose supports are lower sets

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

Here $\text{supp}(\mathbf{z}) := \{\mathbf{i} \in \mathbb{N}_0^d : z_{\mathbf{i}} \neq 0\}$. The corresponding best s -term approximation in lower sets is

$$\sigma_{s,L}(\mathbf{z})_{1,\mathbf{u}} := \inf_{\mathbf{z}' \in \Sigma_{s,L}} \|\mathbf{z} - \mathbf{z}'\|_{1,\mathbf{u}}.$$

Lower sets are common tools in high-dimensional approximation [21, 26, 30, 45, 46]. The lower set structure turns out to be crucial when approximating (a quantity of interest of) the solution map of a parametric PDE. For example, in [22] it has been proved that for a wide class of parametric PDEs the best s -term approximation error (in the mean square and in the uniform sense) in lower sets of the

solution map with respect to Legendre polynomials decays algebraically in s and it is independent of the dimension d of the parametric domain (see also [24, 25]).

An essential property of lower sets is that the union of all lower sets of cardinality s is the well-known *hyperbolic cross* of order s , namely

$$\Lambda_{d,s}^{\text{HC}} := \left\{ \mathbf{i} = (i_1, \dots, i_d) \in \mathbb{N}_0^d : \prod_{\ell=1}^d (i_\ell + 1) \leq s \right\} \equiv \bigcup_{\substack{|S| \leq s \\ S \text{ lower}}} S. \quad (16)$$

From now on, this set will be adopted as our truncated multi-index set, i.e.

$$\Lambda = \Lambda_{d,s}^{\text{HC}}.$$

We note that the cardinality of this index set can be bounded from above as (see, for example, [4])

$$n = |\Lambda_{d,s}^{\text{HC}}| \leq \min \left\{ 2s^3 4^d, e^2 s^{2+\log_2(d)} \right\}. \quad (17)$$

With this in hand, we also define the *intrinsic lower sparsity* of order s as the quantity

$$K(s) := \max \left\{ |S|_{\mathbf{u}} : S \subseteq \mathbb{N}_0^d, |S| \leq s, S \text{ lower} \right\}, \quad (18)$$

where

$$|S|_{\mathbf{u}} := \sum_{\mathbf{i} \in S} u_{\mathbf{i}}^2 \quad (19)$$

is the *weighted cardinality* of a subset S with respect to the weights \mathbf{u} [51]. This quantity is of crucial importance to our results, since it will determine a sufficient condition on the sample complexity m . In the case of Chebyshev and Legendre polynomials, $K(s)$ is known to scale proportionally to s^γ , with

$$\gamma = 2 \quad (\text{Legendre}), \quad \gamma = \ln(3)/\ln(2) \quad (\text{Chebyshev}). \quad (20)$$

Specifically, (see [23, Lemma 3.7]) we have the following:

Lemma 2.2 (Intrinsic lower-sparsity bounds). *Let $2 \leq s \leq 2^{d+1}$. If $\{\phi_{\mathbf{i}}\}_{\mathbf{i} \in \mathbb{N}_0^d}$ is the tensor Legendre or Chebyshev basis, then*

$$s^\gamma/4 \leq K(s) \leq s^\gamma, \quad (21)$$

where $K(s)$ and γ are as in (18) and (20) respectively. Moreover, the upper estimate holds for all $s \geq 2$.

2.4 Robust recovery guarantees under unknown error

We now present the recovery guarantees for the decoders (9)-(12). These estimates are *robust*, in the sense that they admit the presence of unknown error \mathbf{e} corrupting the samples, and *uniform* since, given a set of random pointwise samples distributed according to (6), they hold uniformly in f (with high probability). The proofs of these results are given in §5.

As usual, the expression $X \lesssim Y$ means that there exists a constant C independent of X and Y such that $X \leq CY$. The notation $X \gtrsim Y$ is defined analogously and $X \asymp Y$ means that $X \lesssim Y$ and $X \gtrsim Y$ hold simultaneously.

In all cases (apart from WLAD-LASSO, where an additional constraint appears), uniform recovery with probability at least $1 - \varepsilon$ holds, provided the sample complexity satisfies

$$m \gtrsim s^\gamma \cdot L, \quad (22)$$

where γ is given by (20) and L is a polylogarithmic factor defined as

$$L = L(s, \varepsilon) = \ln^2(s) \min\{\ln(s) + d, \ln(2d) \ln(s)\} + \ln(s) \ln(s/\varepsilon). \quad (23)$$

Note that the particular dependence on d in this estimates stems from bound (17) on the cardinality n of the hyperbolic cross.

We now present the recovery guarantees decoder by decoder.

WQCBP Theorems 5.10 & 5.11 imply that, provided $m \asymp s^\gamma L$, the approximate solution \tilde{f} defined in (13) obtained via WQCBP satisfies

$$\begin{aligned}\|f - \tilde{f}\|_{L^\infty} &\lesssim \sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}} + s^{\gamma/2}[(\eta + \|\mathbf{e}\|_2) + \mathcal{Q}\sqrt{L}\max\{\|\mathbf{e}\|_2 - \eta, 0\}], \\ \|f - \tilde{f}\|_{L_v^2} &\lesssim \frac{\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}}{s^{\gamma/2}} + \eta + \|\mathbf{e}\|_2 + \|f - f_\Lambda\|_{L_v^2} + \mathcal{Q}\sqrt{L}\max\{\|\mathbf{e}\|_2 - \eta, 0\},\end{aligned}$$

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

$$\mathcal{Q} = \mathcal{Q}(\mathbf{A}, \Lambda, \mathbf{u}) := \sqrt{\frac{|\Lambda|\mathbf{u}}{n}} \frac{1}{\sigma_m(\sqrt{\frac{m}{n}}\mathbf{A}^*)},$$

where $\sigma_m(\sqrt{\frac{m}{n}}\mathbf{A}^*)$ is the m^{th} singular value (in decreasing order) of $\sqrt{\frac{m}{n}}\mathbf{A}^*$. Notice that the term depending on $\mathcal{Q}\sqrt{L}$ in the error bound vanishes when $\eta \geq \|\mathbf{e}\|_2$: that is,

$$\eta \geq \|\mathbf{e}\|_2 \quad \Rightarrow \quad \|f - \tilde{f}\|_{L_v^2} \lesssim \frac{\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}}{s^{\gamma/2}} + \eta + \|\mathbf{e}\|_2 + \|f - f_\Lambda\|_{L_v^2},$$

and similarly for the L^∞ error. Moreover, since η and $\|\mathbf{e}\|_2$ both appear in the error bound, the optimal choice of η is

$$\eta = \|\mathbf{e}\|_2. \quad (24)$$

Conversely, when $\|\mathbf{e}\|_2$ is unknown, the price to pay is an additional error term, which depends logarithmically on d and algebraically on s . In particular, the algebraic dependence on s is a consequence of Theorem 5.11, where we show that $\mathcal{Q} \lesssim s^{\alpha/2}/\sigma_m(\sqrt{\frac{m}{n}}\mathbf{A}^*)$, with $\alpha = 1, 2$ for tensor Chebyshev and Legendre polynomials, respectively.

This bound does, however, suggest that tuning η empirically via cross validation, so as to achieve $\eta \approx \|\mathbf{e}\|_2$ can improve the recovery error. Our numerical results in §4 partially support this conclusion.

WLASSO In Theorem 5.13, we show that, provided (22) holds and

$$\lambda \asymp \frac{\sqrt{K(s)}}{\|\mathbf{e}\|_2} \asymp \frac{s^{\gamma/2}}{\|\mathbf{e}\|_2}, \quad (25)$$

the approximate solution \tilde{f} computed by means of WLASSO satisfies

$$\|f - \tilde{f}\|_{L^\infty} \lesssim \sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}} + s^{\gamma/2}\|\mathbf{e}\|_2, \quad (26)$$

$$\|f - \tilde{f}\|_{L_v^2} \lesssim \frac{\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}}{s^{\gamma/2}} + \|\mathbf{e}\|_2 + \|f - f_\Lambda\|_{L_v^2}, \quad (27)$$

with probability at least $1 - \varepsilon$.

Unfortunately, the choice of tuning parameter (25) requires knowledge of $\|\mathbf{e}\|_2$, similar to WQCBP. On the other hand, this requirement is certainly less stringent than the one-sided bound $\eta \geq \|\mathbf{e}\|_2$: no logarithmic factors are present in the error bounds if $\|\mathbf{e}\|_2$ is estimated up to a constant.

WSR-LASSO Theorem 5.14 shows that provided (22) holds and

$$\lambda \asymp \sqrt{K(s)} \asymp s^{\gamma/2}, \quad (28)$$

the approximate solution \tilde{f} computed using WSR-LASSO satisfies (26) and (27) with probability at least $1 - \varepsilon$. In other words, WSR-LASSO attains the same recovery guarantees as WLASSO and WQCBP *but without any prior knowledge* on the noise \mathbf{e} . Specifically, optimal WLASSO tuning parameter (25) depends on $\|\mathbf{e}\|_2$, whereas the optimal WSR-LASSO tuning parameter (28) is independent of this factor. We note in passing that λ depends on the sparsity s , but this is always known in our framework since it is the index of the truncated hyperbolic cross set (16).

WLAD-LASSO From Theorem 5.25, we see that, given $k \in \mathbb{N}$ and provided that

$$m \gtrsim s^\gamma \cdot \max\{L, k\}, \quad (29)$$

the approximate solution \tilde{f} obtained via WLAD-LASSO satisfies

$$\|f - \tilde{f}\|_{L^\infty} \lesssim \sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}} + \lambda \sigma_k(\mathbf{e})_1,$$

and also

$$\|f - \tilde{f}\|_{L_\nu^2} \lesssim (1 + \sqrt{\Theta}) \left(\frac{\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}}{s^{\gamma/2}} + \frac{\sigma_k(\mathbf{e})_1}{\sqrt{k}} \right) + \|f - f_\Lambda\|_{L_\nu^2},$$

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

$$\Theta := \frac{\sqrt{K(s) + \lambda^2 k}}{\min\{\sqrt{K(s)}, \lambda \sqrt{k}\}}, \quad (30)$$

and

$$\sigma_k(\mathbf{e})_1 := \inf_{\mathbf{d}: |\text{supp}(\mathbf{d})| \leq k} \|\mathbf{e} - \mathbf{d}\|_1,$$

where $\text{supp}(\mathbf{d}) := \{i \in [m] : d_i \neq 0\}$, is the best k -term approximation error. Notice that the quantity Θ defined in (30) is minimized for

$$\lambda = \sqrt{\frac{K(s)}{k}}. \quad (31)$$

However, $\lambda = 1$ seems to be a better choice in practice (see Fig. 3 and numerical illustrations in [3]).

Observe that the recovery guarantees for WQCBP, WLASSO and WSR-LASSO all depend on $\|\mathbf{e}\|_2$. Conversely, for WLAD-LASSO this term is replaced by $\sigma_k(\mathbf{e})_1$. This suggests WLAD-LASSO, unlike the other decoders, can effectively correct for errors of type (b). Specifically, if $\mathbf{e} = \mathbf{e}^{\text{bounded}} + \mathbf{e}^{\text{sparse}}$, where $\mathbf{e}^{\text{sparse}}$ has at most k nonzero entries, then

$$\sigma_k(\mathbf{e})_1 \leq \sigma_k(\mathbf{e}^{\text{sparse}})_1 + \sigma_k(\mathbf{e}^{\text{bounded}})_1 \leq \|\mathbf{e}^{\text{bounded}}\|_1 \leq \sqrt{m} \|\mathbf{e}^{\text{bounded}}\|_2.$$

In particular, if $\lambda \asymp s^{\gamma/2}/\sqrt{k}$, as suggested by (31), this yields the error bound

$$\|f - \tilde{f}\|_{L_\nu^2} \lesssim \frac{\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}}{s^{\gamma/2}} + s^{\gamma/2} \|\mathbf{e}^{\text{bounded}}\|_2 + \|f - f_\Lambda\|_{L_\nu^2}, \quad (32)$$

and similarly for the L^∞ -norm error. Up to the factor $s^{\gamma/2}$ (see Remark 2.3 below) this is the same as the bounds for WLASSO and WSR-LASSO. Except, of course, that the WLAD-LASSO corrects for an k -sparse corruption error of arbitrary magnitude, unlike the other decoders.

Remark 2.3 (Sharpness of the sample complexity). The estimate (29) suggests the number of corrupted samples k can be roughly L while retaining the same sample complexity as in the uncorrupted case. We do not believe this is sharp. Indeed, there is reason to expect that $k = cm$ corruptions can be allowed, for some $0 < c < 1$. See [3, 43] for further discussion. If this conjecture were true, then the $s^{\gamma/2}$ dependence in the noise term in (32) would vanish.

3 Literature review

For a general introduction to the optimization programs (9)-(12) in the unweighted case, we refer the reader to, e.g., [31, 61, 37]. We now briefly review the literature regarding each decoder.

WQCBP This optimization program has a long history [28, 44], but here we are particularly interested in its application in compressed sensing [18, 27, 31]. The weighted version of QCBP has been introduced in the context of function approximation with bounded orthonormal systems in [51], although previously studied in [19, 32, 48, 68]. For applications of WQCBP to high-dimensional function approximation and uncertainty quantification, we refer the reader to [1, 2, 4, 23, 48, 50, 67] and references therein. All these works focus on the error-aware scenario, where upper bounds of the form (1) are assumed to be known *a priori*, even though this assumption is unlikely to be met in practice. It is worth mentioning that robust recovery guarantees for unweighted ℓ^1 minimization have been recently proved in the error-blind scenario, where upper bounds of the form (1) are not assumed to be known (see [15, 16] and references therein). In this paper, these results are generalized to the weighted case.

WLASSO The literature regarding LASSO (Least Absolute Shrinkage and Selection Operator) is boundless. The pioneering paper [60] by R. Tibshirani has already reached more than twenty thousand citations according to Google Scholar (roughly one thousand citations per year on average). We refer the reader to [37, §2.10] and references therein for an extensive review on the state-of-the-art of LASSO and for historical remarks. A weighted variant of LASSO, called the “adaptive LASSO” due to the iterative and adaptive procedure used to update the weights, is considered in [39, 69].

WSR-LASSO The unweighted version of SR-LASSO has been introduced in [12] and studied under the name “scaled LASSO” in [58]. Further results about SR-LASSO can be found in [8, 13, 17, 49, 59]. In [54], the authors consider a version of the SR-LASSO where the ℓ^1 norm is replaced by a generic sparsity-inducing norm, admitting the weighted SR-LASSO as a particular case. For the statisticians, an attractive feature of SR-LASSO is that the optimal tuning parameter does not depend on the variance of the random noise corrupting the observations, like in the LASSO case. Our study confirms that this is the case also in high-dimensional function approximation.

WLAD-LASSO The LAD (Least Absolute Deviation) regression and the LASSO program have been combined into the so-called LAD-LASSO decoder in order to make the LASSO more robust to the presence of heavy-tailed errors or outliers in the response [33, 34, 62, 64, 65]. In the context of compressed sensing, this decoder is also known as a *fault-tolerant* version of ℓ^1 minimization. It has been considered in [41, 43, 47, 53, 55, 56, 57, 63] and in [3, 52] with applications to uncertainty quantification. In [7], the author introduces the weighted variant of LAD-LASSO and proves that it has some statistical properties (asymptotic normality and consistency). An adaptively reweighted version of LAD-LASSO is considered in [42].

4 Numerical experiments

We compare the decoders (9)-(12) from the numerical viewpoint. The section is structured as follows. First, we formally define the cross-validation procedure employed to choose the tuning parameters in §4.1. Then, we compare the four decoders by studying their recovery error as a function of the tuning parameter (§4.2) and of the sample complexity (§4.3). Finally, we compare the performance of the four decoders in an application to the uncertainty quantification of parametric ODEs and PDEs with random inputs in §4.4.

All the numerical experiments have been performed in MATLAB[®] using CVX, an optimization toolbox for solving convex problems [35, 36]. We always run CVX setting `cvx_precision best` and `cvx_solver mosek`. We have used MATLAB[®] R2016b version 9.1 64-bit on a MacBook Pro equipped with a 3 GHz Intel Core i7 processor and with 8 GB DDR3 RAM. For the sake of convenience, we will sometimes use the MATLAB[®] vector notation to denote objects like $10.^{(1 : 0.5 : 2)} = (10^1, 10^{1.5}, 10^2)$.

4.1 Cross validation

We define the cross-validation procedure that will be employed in the next sections. The pseudo-code is reported in Algorithm 4.1. It corresponds to the cross-validation procedure described in [37, §2.3], usually referred to as *K-fold cross validation* [6]. It is worth noticing that step 9 of Algorithm 4.1 can be replaced with $\varepsilon(t, g, p) = \|\mathbf{A}_v \hat{\mathbf{x}} - \mathbf{y}_v\|_2$, as it is done in [29, 67]. However, we prefer to use the squared residual, in accordance with [37]. In addition to the design matrix \mathbf{A} and the vector of samples \mathbf{y} , the cross-validation procedure takes as input:

- a number $G \in \mathbb{N}$ of sample groups;
- a decoder Δ with tuning parameter p such that $\hat{\mathbf{x}} = \Delta(\mathbf{A}, \mathbf{y}; p)$ approximately solves $\mathbf{A}\mathbf{x} = \mathbf{y}$;
- a finite set \mathcal{P} of parameters;
- a number $T \in \mathbb{N}$ of repeated random tests to run.

The procedure gives as output a tuning parameter $p_{\text{cv}} \in \mathcal{P}$.

Algorithm 4.1 Cross-validation procedure

```

1: procedure  $p_{\text{cv}} = \text{CROSSVALIDATION}(\mathbf{A}, \mathbf{y}, G, \Delta, \mathcal{P}, T)$ ;
2: for all  $t \in [T]$  do
3:   Randomly partition  $[m] = I_1 \sqcup \dots \sqcup I_G$  into  $G$  sets of cardinality  $\lfloor m/G \rfloor$  or  $\lfloor m/G \rfloor + 1$ ;
4:   for all  $g \in [G]$  do
5:      $\mathbf{A}_v = \sqrt{\frac{m}{|I_g|}} (A_{ij})_{i \in I_g, j \in [n]}, \quad \mathbf{y}_v = \sqrt{\frac{m}{|I_g|}} (y_i)_{i \in I_g};$ 
6:      $\mathbf{A}_r = \sqrt{\frac{m}{m-|I_g|}} (A_{ij})_{i \in [m] \setminus I_g, j \in [n]}, \quad \mathbf{y}_r = \sqrt{\frac{m}{m-|I_g|}} (y_i)_{i \in [m] \setminus I_g};$ 
7:     for all  $p \in \mathcal{P}$  do
8:        $\hat{\mathbf{x}} = \Delta(\mathbf{A}_r, \mathbf{y}_r; p)$ ;
9:        $\varepsilon(t, g, p) = \|\mathbf{A}_v \hat{\mathbf{x}} - \mathbf{y}_v\|_2^2$ ;
10:    end for
11:  end for
12: end for
13:  $p_{\text{cv}} = \arg \min_{p \in \mathcal{P}} \frac{1}{T \cdot G} \sum_{t \in [T]} \sum_{g \in [G]} \varepsilon(t, g, p).$ 

```

4.2 Tuning parameter *vs.* error

We compare WQCBP, WLASSO, WSR-LASSO, and WLAD-LASSO by studying the behavior of the recovery error as a function of the tuning parameter when approximating the function

$$f(t) = \exp \left(-\frac{1}{d} \sum_{\ell=1}^d \cos(t_\ell) \right), \quad \text{with } d = 15. \quad (33)$$

The aim of this experiment is to both compare the performance of the four decoders, and to validate the optimal choices of the tuning parameters suggested by the theory (see §2.4).

We consider a sparsity level $s = 10$, corresponding to $n = |\Lambda_{15,10}^{\text{HC}}| = 1432$ and a number of samples $m = \lceil s^\gamma \log(n) \rceil$, with γ defined as in (20). In particular, this corresponds to a sample complexity $m = 727$ for Legendre and $m = 280$ for Chebyshev polynomials. Then, we repeat the following experiment 50 times: we generate m random samples and corrupt them by random noise $\mathbf{e} = \beta \mathbf{n} / \|\mathbf{n}\|_2$ with $\beta \in \{0, 10^{-3}, 10^{-2}, 10^{-1}\}$, where $\mathbf{n} \in \mathbb{R}^m$ is a random vector with independent entries uniformly distributed over $[-1, 1]$, and we solve the resulting system (7) by means of WQCBP, WLASSO, WSR-LASSO, and WLAD-LASSO for each value of the tuning parameter, as specified in Table 1.²

²In practice, for WLAD-LASSO we use $\lambda = 1.01$ instead of $\lambda = 1$ since the choice $\lambda = 1$ leads to the presence of spurious outliers in the box plot. We think that this behavior is due to CVX and not to the decoder itself.

WQCBP (η)	WLASSO (λ)	WSR-LASSO (λ)	WLAD-LASSO (λ)
$10.^{-7:0.5:1}$	$10.^{-1:0.5:8}$	$10.^{-2:0.25:5}$	$10.^{-2:0.25:3}$

Table 1: Sets of tuning parameters for WQCBP, WLASSO, WSR-LASSO, and WLAD-LASSO used to generate Figs. 1 & 2.

The results are shown in Figs. 1 and 2. In order to make the statistics of the randomized experiments transparent, we visualize the results in a box plot. The continuous lines represent the median values (in accordance with the box plot information). We consider the L_ν^2 error with respect to a high-fidelity approximation to f , computed using least squares and $20n = 28620$ random pointwise evaluations.³

It is remarkable that all the decoders are able to reach an accuracy below the noise level for suitable values of the parameter. Moreover, the optimal choice of the tuning parameters seems to be in accordance with the theory provided in §2.4. Indeed, we observe a direct proportionality between the optimal value of η and the noise level for WQCBP, and an inverse proportionality between the optimal value of λ and the noise level for WLASSO. For WSR-LASSO and WLAD-LASSO, the optimal values of λ are independent on the noise level. We also notice that the limit value of the L_ν^2 recovery error associated with the unconstrained optimization programs for $\lambda \rightarrow \infty$ coincides with the limit value of the L_ν^2 error associated with WQCBP for $\eta \rightarrow 0$. This suggests that the solution to the unconstrained programs tends to the solution of WBP as $\lambda \rightarrow \infty$. In other words, choosing a very large λ forces the data fidelity constraint to be realized exactly, as it is natural to expect.

Let us take a closer look to the performance of WSR-LASSO and WLAD-LASSO, recalling the optimal choices for λ given by (28) and (31), respectively. For WSR-LASSO, being $s = 10$, we have $\sqrt{K(s)} \approx 10$ for Legendre and $\sqrt{K(s)} \approx \sqrt{10^{\log(3)/(2\log(2))}} \approx 6.2$ for Chebyshev polynomials. From the box plots, the optimal value of λ seems to be around $10^{1.5} \approx 31.6$ (Legendre) and $10^{1.25} \approx 17.8$ (Chebyshev). Therefore, $\lambda \approx 3\sqrt{K(s)}$. For the WLAD-LASSO, in principle the theory predicts that $\lambda \approx \sqrt{K(s)/K}$. Since the error e is nonsparse in this case, we have $K = m \approx K(s)\log(n)$. We see that $\lambda \approx 3/\sqrt{\log(n)} \approx 1.1129$ seems to match the numerics well. This is the reason why we choose the hidden constant to be 3 in (28) and (31) for the subsequent experiments.

4.3 Sampling complexity *vs.* error

We make another comparison of the four decoders by studying the decay of the L_ν^2 error as a function of the sampling complexity m when approximating the function f defined in (33) with $d = 10$, employing tensorized Chebyshev polynomials. We set $s = 15$, corresponding to $n = |\Lambda_{10,15}^{\text{HC}}| = 1341$ and consider a number of samples $m = \lceil C \cdot K(s) \rceil$ for $C = 2 : 0.5 : 4$, where $K(s) \approx 15^{\log(3)/\log(2)} \approx 73.1$. First, we compute a high-fidelity solution $\hat{\mathbf{x}}_{\text{LS}}$ by means of least squares, using $20 \cdot n = 26820$ random pointwise samples. Then, we consider ten different combinations of solvers and parameters. For the sake of the comparison, we use the high-fidelity solution $\hat{\mathbf{x}}_{\text{LS}}$ to produce “oracle” estimates of the noise level (notice that this is an idealized scenario since $\hat{\mathbf{x}}_{\text{LS}}$ is not available in practice). Moreover, we employ the cross-validation procedure described in Algorithm 4.1, and the values of the parameters suggested by the theoretical analysis for WSR-LASSO ($\lambda = 3\sqrt{K(s)}$) and WLAD-LASSO ($\lambda = 3/\sqrt{\log(n)}$) (see §2.4 and the discussion at the end of §4.2). The cross-validation procedure is always employed as $\text{CROSSVALIDATION}(\mathbf{A}, \mathbf{y}, 5, \Delta, \mathcal{P}, 3)$, where Δ and \mathcal{P} are specified on a case-by-case basis. The ten combinations considered are the following:

1. WQCBP with $\eta = 0$;
2. WQCBP with $\eta = \eta_{\text{oracle}} = \|\mathbf{A}\hat{\mathbf{x}}_{\text{LS}} - \mathbf{y}\|_2$;
3. WQCBP with cross validation, where $\mathcal{P} = \eta_{\text{oracle}} \cdot 10.^{-2:0.5:2}$;
4. WLASSO with $\lambda = \lambda_{\text{oracle}} = \sqrt{K(s)}/\|\mathbf{A}\hat{\mathbf{x}}_{\text{LS}} - \mathbf{y}\|_2$;

³Notice that the outliers are sometimes aligned (e.g., in the tail of the blue curve in Fig. 1 bottom right). This is due to the structure of the proposed numerical experiment: for each randomized choice of samples, all the parameters are tested using the same samples.




Figure 1: Box plot of the recover L^2_ν error as a function of the tuning parameter for WQCBP (top left), Wlasso (top right), WSR-LASSO (bottom left), and WLAD-LASSO (bottom right) when computing the sparse approximation of the function f defined in (33) with respect to tensorized Legendre polynomials.

5. Wlasso with cross validation, where $\mathcal{P} = \lambda_{\text{oracle}} \cdot 10.^{(-2:0.5:2)}$;
6. WSR-LASSO with $\lambda = 3\sqrt{K(s)}$;
7. WSR-LASSO with , where $\mathcal{P} = 3\sqrt{K(s)} \cdot 10.^{(-2:0.5:2)}$;
8. WLAD-LASSO with $\lambda = 3/\sqrt{(k/m)\log(n)}$;
9. WLAD-LASSO with cross validation, where $\mathcal{P} = \frac{3\sqrt{H}}{\sqrt{m\log(n)}} \cdot 10.^{(-2:0.5:2)}$;
10. WLAD-LASSO with $\lambda = 1$.

Notice that for WLAD-LASSO we set $\mathbf{v} = 1$ in (12).

This comparison is performed in three different scenarios. First, without corrupting the samples, thus having only truncation error (8) (Fig. 3, top right). Second, corrupting the measurement by random noise $\mathbf{e} = 10^{-2}\mathbf{n}/\|\mathbf{n}\|_2$, where \mathbf{n} is a random vector with independent entries uniformly distributed over $[-1, 1]$ (Fig. 3, bottom left). Third, corrupting 10% of the measurements by a random noise uniformly distributed over $[-10, 10]$ (Fig. 3, bottom right). The respective curves show the L^2_ν error with respect to the high-fidelity solution, averaged over 25 trials.

Let us comment the results in Fig. 3. First, the choices of the parameters predicted by the theory give good results in general. Moreover, the performance of each decoder is usually improved by cross validation. WBP slightly underperforms with respect to all the other solvers apart from the case without extra noise and for m large enough (Fig. 3 top right). The performance of WSR-LASSO is comparable with those of WQCBP and Wlasso, but for WSR-LASSO the optimal parameter choice does not

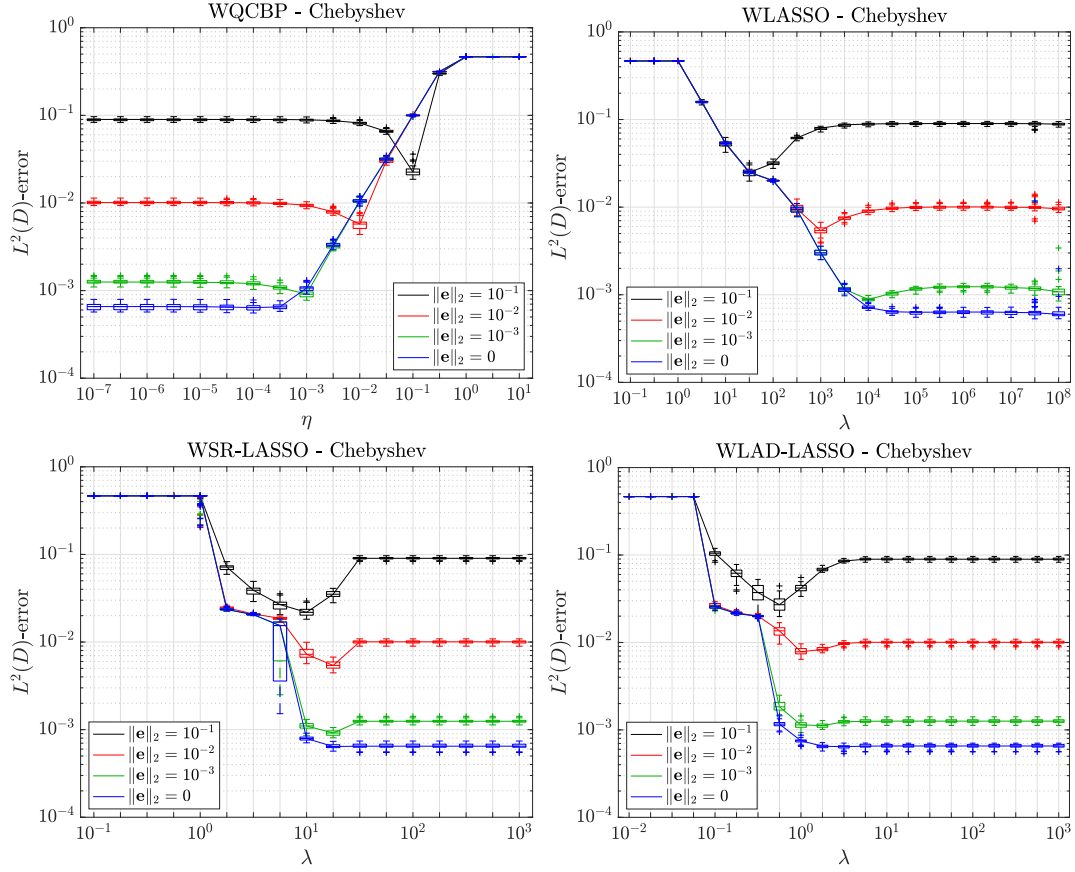


Figure 2: Same experiment as in Fig. 1, but for tensorized Chebyshev polynomials.

depend on the noise level. Finally, in the case of sparsely corrupted measurements, WLAD-LASSO outperforms the other decoders, as expected (Fig. 3 bottom right). The choice $\lambda = 1$ is particularly good in the sparsely corrupted case. It is also worth noting that cross validation does not seem to help in this case. This is natural, since Algorithm 4.1 chooses the tuning parameter that minimizes the ℓ^2 norm of the residual, which is not small in this case. A possible remedy could be to replace step 9 in Algorithm 4.1 with $\varepsilon(t, g, p) = \|\mathbf{A}_v \hat{\mathbf{x}} - \mathbf{y}_v\|_1$.

We note in passing that the experiments in Figs. 1, 2 & 3 yield analogous results when the random noise \mathbf{n} used to corrupt the samples has independent normal Gaussian entries $\mathcal{N}(0, 1)$. We do not add the resulting figures for the sake of brevity.

4.4 Application to parametric ODEs and PDEs with random inputs

Finally, we compare the performance of the decoders under exam when the function begin approximated is a quantity of interest of the solution to a parametric ODE or PDE random inputs.

4.4.1 Parametric PDEs: Diffusion equation

We consider the following diffusion equation:

$$\begin{cases} -\nabla \cdot (a_t(x, y) \nabla u_t(x, y)) = F, & \forall (x, y) \in \Omega, \\ u_t(x, y) = 0, & \forall (x, y) \in \partial\Omega, \end{cases} \quad (34)$$

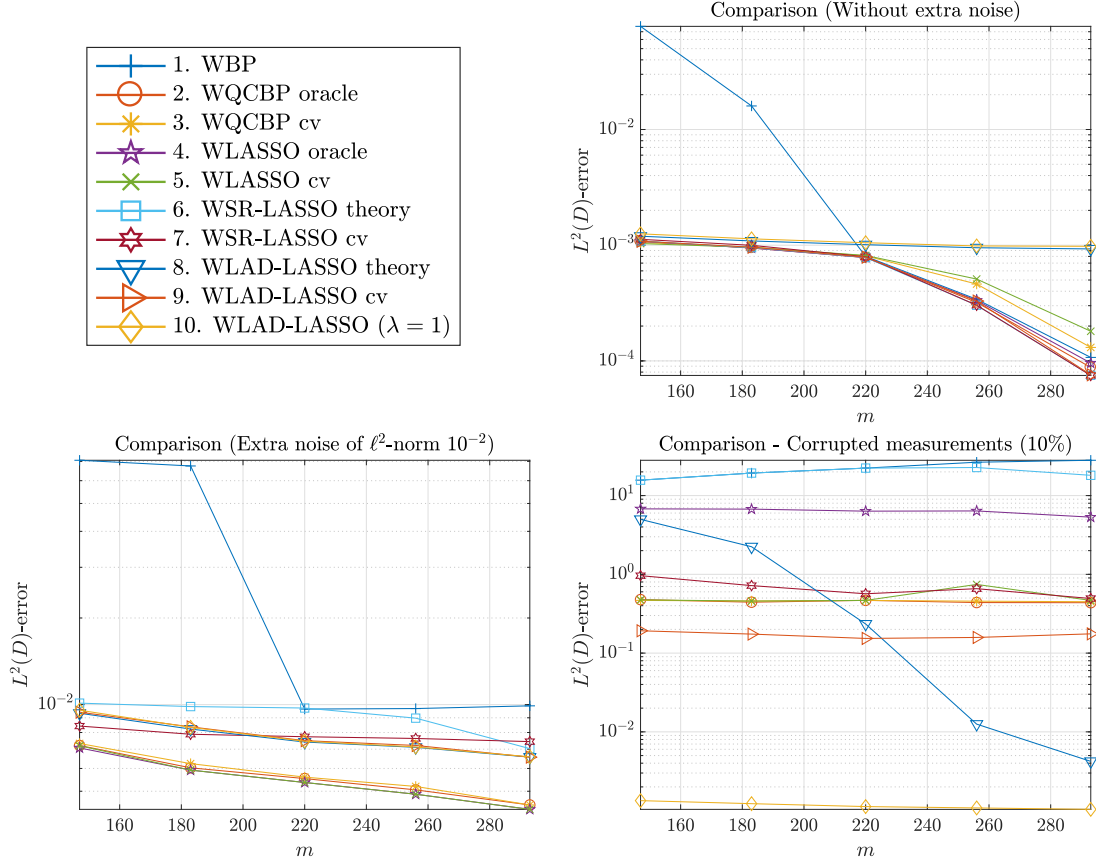


Figure 3: Comparison among WQCBP, WLASSO, WSR-LASSO, and WLAD-LASSO for the approximation of (33) with $d = 10$. The ten combinations of decoder and parameter considered are described in §4.3.

where $\Omega = [0, 1]^2$ is the physical domain, a_t is a parametric diffusion coefficient, u_t is the unknown solution, and F is a nonparametric forcing term. This example and analogous variations have been considered in [9, 10, 20]. The parametric space is $D = (-1, 1)^8$ and the diffusion coefficient affinely depends on the parameters

$$a_t(x, y) = 1 - \sum_{\ell=1}^8 \mathbb{1}_{\Omega_\ell}(x, y)(0.595 + 0.395t_\ell), \quad \forall (x, y) \in \Omega, \forall t \in D,$$

where $\Omega_1, \dots, \Omega_8$ are circular subregions of Ω of radius 0.13 placed symmetrically with respect to the center of Ω (see Fig. 4). If we think about t as a random vector uniformly distributed in D , then $a_t|_{\Omega_\ell}$ can be thought as a random variable uniformly distributed $[0.01, 0.99]$. The forcing term is

$$F(x, y) = 100 \cdot \mathbb{1}_{\Omega_F}(x, y), \quad \forall (x, y) \in \Omega,$$

where $\Omega_F = [0.4, 0.6]^2$. We are interested in the function defined by the following quantity of interest:

$$f(t) = \int_{\Omega_F} u_t(x, y) dx dy. \quad (35)$$

We sample $f(t)$ by solving the PDE (34) with FreeFem++ (version 3.42, 64 bits), employing triangular P1 elements and generating finer and finer meshes [38]. The mesh is generated with the FreeFem++ mesher, by controlling its resolution by a parameter N_{el} . In particular, we parametrize each edge of Ω

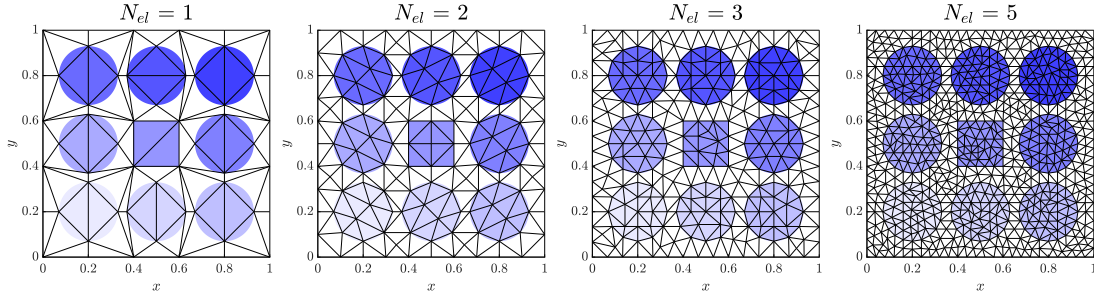


Figure 4: The physical domain Ω discretized with increasing mesh resolution, corresponding to $N_{el} = 1, 2, 3, 5$ (from left to right). The circular regions $\Omega_1, \dots, \Omega_8$ and the central square region Ω_F are highlighted.

using $5N_{el}$ elements, the circumference of each Ω_ℓ with $4N_{el}$ elements and each edge of Ω_R using N_{el} elements. Then, we let $N_{el} = 1, 2, 3, 5$. The corresponding meshes are represented in Fig. 4 and have a number of degrees of freedom equal to 56, 141, 267, and 715, respectively. Each sample $f(\mathbf{t}_i)$ is then computed by considering the integral of the resulting piecewise linear approximation to $u_{\mathbf{t}_i}$ over Ω_F . In this case, the unknown error affecting the samples contains the truncation error (8), the discretization error due to the Galerkin projection, and the numerical error associated with the FreeFem++ solver. We only expect this unknown error to decrease with respect to N_{el} .

Similarly to §4.2 & 4.3, we study the behavior of the recovery error as a function of the tuning parameter and of the sample complexity. Since \mathbf{t} is uniformly distributed in D , we employ tensorized Legendre polynomials.

In Fig. 5 we consider an experiment analogous to Figs. 1 & 2 for the function f defined in (35). Namely, we plot the L_ν^2 error as a function of the tuning parameter (chosen as in Table 1) for the four decoders. We set $s = 10$, resulting in $n = 353$, and $m = 30$. For each value of $N_{el} = 1, 2, 3, 5$, we run 50 random experiments. The L_ν^2 error is computed with respect to a high-fidelity solution computed using a mesh size $N_{el} = 100$ (with 267076 degrees of freedom) and approximated with least squares from 50000 random samples.

As expected, the best accuracy achieved by each solver is a decreasing function of N_{el} . We can still see that the best choice of the parameter depends on the noise level for WQCBP and WLASSO, whereas it is independent of the noise for WSR-LASSO and WLAD-LASSO. Surprisingly, the strict global minima corresponding to the optimal choice of the parameter observed in Figs. 1 & 2 do not appear here. In fact, all solvers exhibit a plateau. We may argue that the optimal choices (24), (25), (28), and (31) hold under the form of inequality, and not of equality. Namely, $\eta \lesssim \|e\|_2$ for WQCBP, $\lambda \gtrsim \sqrt{K(s)}/\|e\|_2$ for WLASSO, $\lambda \gtrsim \sqrt{K(s)}$ for WSR-LASSO, and $\lambda \gtrsim 1/\sqrt{\log(n)}$ for WLAD-LASSO.⁴ See §4.4.3 for some further discussion.

Next, we compare the performance of the four decoders using the ten combinations of decoders and tuning parameters described in §4.3. We let $m = (20 : 20 : 100)$ and use $N_{el} = 5$. Fig. 6 shows the L_ν^2 error as a function of m , averaged over 25 trials. Overall, the decoders have similar performances, in that the recovery error differs by less than one digit. Cross validation is able to achieve the best accuracy level in all cases. Remarkably, WSR-LASSO achieves the best accuracy using the choice (28) suggested by the theory.

⁴We have performed the same experiment for quantities of interest different from (35), such as the integral of $u_{\mathbf{t}}$ over the regions Ω_i , pointwise evaluations of $u_{\mathbf{t}}$, or the integral of $u_{\mathbf{t}}^2$ over Ω_i or Ω_F . In all these cases, we do not observe the strict global minima in the parameter-vs-error plot. These experiments are not reported here for the sake of brevity.

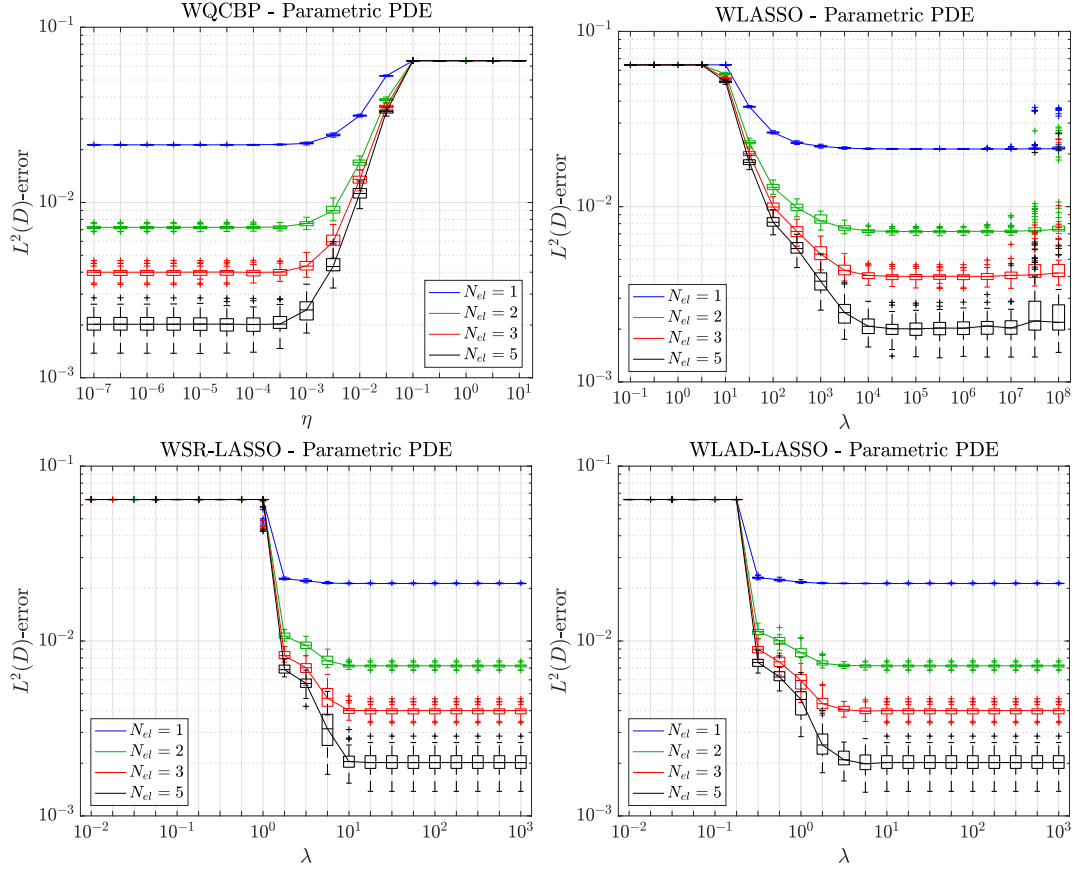


Figure 5: Box plot of the recovery error as a function of the tuning parameter for WQCBP (top left), WLASSO (top right), WSR-LASSO (bottom left), and WLAD-LASSO (bottom right) for the sparse approximation of the quantity of interest f defined in (35) with respect to tensorized Legendre polynomials.




Figure 6: Comparison among WQCBP, WLASSO, WSR-LASSO, and WLAD-LASSO for the approximation of the quantity of interest f defined in (35).

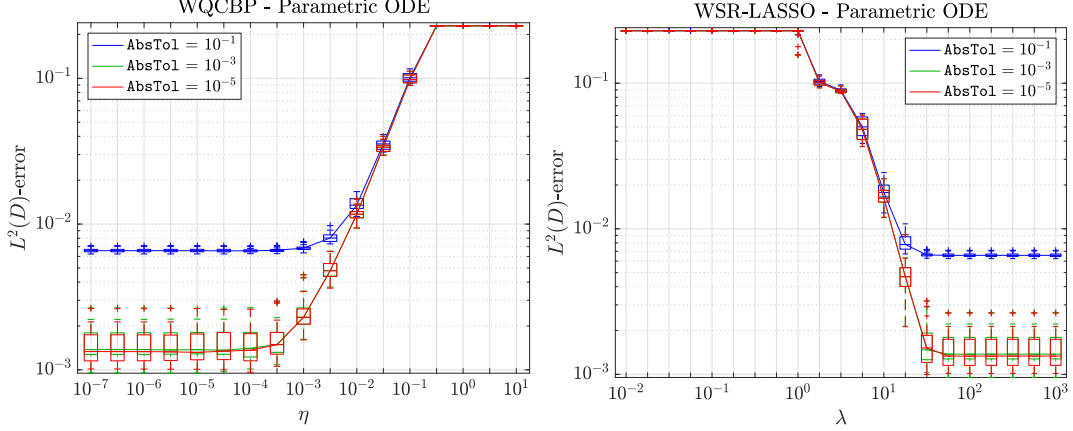


Figure 7: Boxplot of the L^2_ν error as a function of the tuning parameter for WQCBP (left) and WSR-LASSO (right) when approximating the quantity of interest (37) in the damped harmonic oscillator case.

4.4.2 Parametric ODEs: Damped harmonic oscillator

We study an application to parametric ODEs with random inputs, considering the equation of a damped harmonic oscillator subject to external forcing

$$\begin{cases} u_t''(x) + \gamma u_t'(x) + k u_t(x) = g \cos(\omega x), & \forall x > 0, \\ u_t(0) = u_0, & u_t'(0) = v_0, \end{cases} \quad (36)$$

where the parameters $\gamma, k, g, \omega, u_0, v_0$ are defined as a function of a random vector \mathbf{t} uniformly distributed in $D = [-1, 1]^6$ as follows:

$$\begin{aligned} \gamma &= 0.1 + 0.02t_1 \in [0.08, 0.12], & k &= 0.035 + 0.05t_2 \in [0.03, 0.04], & g &= 0.1 + 0.02t_3 \in [0.08, 0.12], \\ \omega &= 1 + 0.2t_4 \in [0.8, 1.2], & u_0 &= 0.5 + 0.05t_5 \in [0.45, 0.55], & v_0 &= 0.05t_6 \in [-0.05, 0.05]. \end{aligned}$$

These restrictions make the oscillator always underdamped. We are interested in approximating the following quantity of interest:

$$f(\mathbf{t}) = u_{\mathbf{t}}(20). \quad (37)$$

An analogous experiment is considered in [3].

We generate the pointwise samples by solving (36) in MATLAB[®] using the command `ode23` and setting different values for the tolerance parameter `AbsTol` in $\{10^{-1}, 10^{-3}, 10^{-5}\}$ within `odeset`. We employ the Legendre polynomials as sparsity basis and set $s = 20$ (corresponding to $n = 795$) and $m = 100$. In Fig. 7 (left), we show the boxplot of the L^2_ν error as a function of the parameter η for WQCBP and WSR-LASSO over 25 runs, where the L^2_ν error is computed with respect to a high-fidelity reference solution \mathbf{x}_{ref} , approximated via least squares from 10000 random samples with `AbsTol` = 10^{-14} .

Analogously to the case of the parametric PDE (Fig. 5 top left), we see that there are no strict global minima. The optimal parameter choices appear to be $\eta \lesssim \|e\|_2$ and $\lambda \gtrsim \sqrt{K(s)}$ as in the previous experiment.⁵

4.4.3 Discussion

Both the parametric PDE and ODE problems considered fail to exhibit the strict global minima in the error versus parameter plots as seen in the synthetic examples (Figs. 1 & 2). We conjecture that this stems from the fact that errors produced by the numerical solvers are structured, as opposed to

⁵The same phenomenon is observed for WLASSO and WLAD-LASSO, but the plots are not shown here for the sake of brevity.

random. Note that our theoretical recovery guarantees are adversarial, i.e. uniform with respect to the 2-norm $\|\mathbf{e}\|_2$ of the error $\mathbf{e} \in \mathbb{C}^m$, and therefore cannot explain this behaviour. Evidently, the structured errors produced by the solvers are quite far from this worst-case scenario. This may imply that careful parameter tuning is not required in practice, although such a conclusion requires more broad examination. In particular, understanding this phenomenon from the theoretical viewpoint remains an open problem.

5 Theoretical analysis

This section contains the proofs of the results of §2.4. It is organized in two main parts. First, we prove recovery error estimates for WQCBP, WLASSO, and WSR-LASSO in §5.1. Then, in §5.2 we deal with WLAD-LASSO. This section structure depends on the fact that the proof strategy adopted for the WLAD-LASSO contains some additional technicalities with respect to the other three decoders.

Remark 5.1 (Proof strategy). In order to prove robust recovery error estimates, we will follow a conceptual roadmap commonly employed in compressed sensing [31]. This is composed of the following elements: (LB) lower bound on the sample complexity m ; (RIP) restricted isometry property for \mathbf{A} ; (RNSP) robust null space property for \mathbf{A} ; (DB) distance bound; (REE) recovery error estimate. The chain of implications showed is then:

$$(\text{LB}) \Rightarrow (\text{RIP})^{\text{whp}} \Rightarrow (\text{RNSP})^{\text{whp}} \Rightarrow (\text{DB})^{\text{whp}} \Rightarrow (\text{REE})^{\text{whp}},$$

where the notation $(x)^{\text{whp}}$ indicates that the property (x) holds with high probability.

5.1 WQCBP, WLASSO, and WSR-LASSO

Before showing recovery error estimates for WQCBP, WLASSO, and WSR-LASSO, we need to introduce suitable versions of the restricted isometry property and of the robust null space property, adapted to the lower set setting (§5.1.1). Then, we prove the recovery error estimates for WQCBP, WLASSO, and WSR-LASSO in §5.1.2, 5.1.3, and 5.1.4, respectively.

5.1.1 Preliminaries

We first recall the lower restricted isometry property and of lower robust null space property [23]:

Definition 5.2 (Lower restricted isometry property). A matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ is said to have the *lower restricted isometry property* of order s if there exists a constant $0 < \delta < 1$ such that

$$(1 - \delta)\|\mathbf{z}\|_2^2 \leq \|\mathbf{A}\mathbf{z}\|_2^2 \leq (1 + \delta)\|\mathbf{z}\|_2^2, \quad \forall \mathbf{z} \in \mathbb{C}^n, \quad |\text{supp}(\mathbf{z})|_{\mathbf{u}} \leq K(s),$$

where $\text{supp}(\mathbf{z}) := \{i \in [n] : z_i \neq 0\}$ and $|\text{supp}(\mathbf{z})|_{\mathbf{u}}$ is its weighted cardinality defined as in (19). The smallest constant such that this property holds is called the s^{th} lower restricted isometry constant of \mathbf{A} and it is denoted as $\delta_{s,L}$.

Definition 5.3 (Lower robust null space property). Given $0 < \rho < 1$ and $\tau > 0$, a matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ is said to have the *lower robust null space property* of order s if

$$\|\mathbf{z}_S\|_2 \leq \frac{\rho}{\sqrt{K(s)}} \|\mathbf{z}_{S^c}\|_{1,\mathbf{u}} + \tau \|\mathbf{A}\mathbf{z}\|_2, \quad \forall \mathbf{z} \in \mathbb{C}^n,$$

for any $S \subseteq \Lambda$ such that $|S|_{\mathbf{u}} \leq K(s)$, where $K(s)$ is defined as in (18).

The lower restricted isometry and the lower robust null space properties of order s are equivalent to the weighted restricted isometry and weighted robust null space properties of order $K(s)$, respectively (see [51, Definition 1.3] and [51, Definition 4.1]).

The following result, corresponding to [23, Proposition 4.4], states that the lower restricted isometry property implies the lower robust null space property up to a suitable restriction on the lower restricted isometry constant.

Lemma 5.4 (Lower restricted isometry property \Rightarrow lower robust null space property). *Let $s \geq 2$ and $\mathbf{A} \in \mathbb{C}^{m \times n}$ satisfy the lower restricted isometry property of order αs with constant*

$$\delta = \delta_{\alpha s, L} < \frac{1}{5},$$

where $\alpha = 2$ if the intrinsic weights \mathbf{u} arise from the tensor Legendre basis and $\alpha = 3$ if the weights arise from the tensor Chebyshev basis. Then \mathbf{A} has the lower robust null space property of order s with

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

In [23, Theorem 2.2] it is proved that a sufficient condition for the lower restricted isometry property of order s to hold with high probability is that the sample complexity m must scale linearly with the intrinsic sparsity $K(s)$ up to a polylogarithmic factor. This result is stated below.

Theorem 5.5 (Sample complexity \Rightarrow lower restricted isometry property). *Fix $0 < \varepsilon < 1$, $0 < \delta < 1$, let $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ be as in (2) and \mathbf{u} be the intrinsic weights defined in (14) and suppose that*

$$m \gtrsim K(s) \cdot L(s, n, \delta, \varepsilon), \quad (38)$$

where $K(s)$ is as in (18) and where $L = L(s, n, \delta, \varepsilon)$ is a polylogarithmic factor defined as

$$L = \frac{1}{\delta^2} \ln \left(\frac{K(s)}{\delta^2} \right) \max \left\{ \frac{1}{\delta^4} \ln \left(\frac{K(s)}{\delta^2} \ln \left(\frac{K(s)}{\delta^2} \right) \right) \ln(n), \frac{1}{\delta} \ln \left(\frac{1}{\delta \varepsilon} \ln \left(\frac{K(s)}{\delta^2} \right) \right) \right\}. \quad (39)$$

Then with probability at least $1 - \varepsilon$ the design matrix \mathbf{A} defined in (5) satisfies the lower restricted isometry property of order s with constant $\delta_{s, L} \leq \delta$.

A fundamental feature of the lower (or weighted) robust null space property is that it implies a suitable $\ell_{\mathbf{u}}^1$ distance bound. This is stated in the following theorem, an immediate consequence of [51, Theorem 4.2] (notice that $\sigma_{K(s)}(\mathbf{z})_{1, \mathbf{u}} \leq \sigma_{s, L}(\mathbf{z})_{1, \mathbf{u}}$, for any $\mathbf{z} \in \mathbb{C}^n$). This can also be viewed as a corollary of the Theorem 5.18, which deals with the framework of 2-level sparsity.

Theorem 5.6 (Lower robust null space property $\Rightarrow \ell_{\mathbf{u}}^1$ distance bound). *If the matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ has the lower null space property of order s , then for every $\mathbf{z}, \hat{\mathbf{z}} \in \mathbb{C}^n$ the following holds:*

$$\|\mathbf{z} - \hat{\mathbf{z}}\|_{1, \mathbf{u}} \leq \frac{1+\rho}{1-\rho} (2\sigma_{s, L}(\mathbf{z})_{1, \mathbf{u}} + \|\hat{\mathbf{z}}\|_{1, \mathbf{u}} - \|\mathbf{z}\|_{1, \mathbf{u}}) + \frac{2\tau\sqrt{K(s)}}{1-\rho} \|\mathbf{A}(\hat{\mathbf{z}} - \mathbf{z})\|_2.$$

The next step is to convert an $\ell_{\mathbf{u}}^1$ error estimate into an L_{ν}^2 error estimate. In order to do this, we need to recall two technical lemmas. Namely, an upper bound on the ℓ^∞ norm of the intrinsic weights and a weighted version of the Stechkin inequality. Their proofs can be found in [23, Lemma 4.1] and [51, Theorem 3.2], respectively.

Lemma 5.7 (ℓ^∞ bound on the intrinsic weights). *For any $s \geq 2$ and $d \geq 1$, let $\Lambda = \Lambda_{d, s}^{\text{HC}}$. Then*

$$\|\mathbf{u}_\Lambda\|_\infty^2 \leq \frac{3}{4} K(s),$$

where $(u_i)_{i \in \mathbb{N}_0^d}$ are the weights (14) associated with the tensor Chebyshev or Legendre polynomials.

Lemma 5.8 (Weighted Stechkin inequality). *For any $\mathbf{z} \in \ell_{\mathbf{u}}^1(\mathbb{N}_0^d)$ and $K \in \mathbb{N}$, with $K > \|\mathbf{u}\|_\infty^2$, the following holds*

$$\inf \{ \|\mathbf{z} - \mathbf{z}'\|_2 : |\text{supp}(\mathbf{z}')|_{\mathbf{u}} \leq K \} \leq \frac{\|\mathbf{z}\|_{1, \mathbf{u}}}{\sqrt{K - \|\mathbf{u}\|_\infty^2}}.$$

Thanks to Lemma 5.7, Lemma 5.8, we are able to give a way to convert an $\ell_{\mathbf{u}}^1$ error estimate (on the coefficients) into an L_{ν}^2 error estimate (on the function) under the lower robust null space property.

Lemma 5.9 (Lower robust null space property $\Rightarrow \ell_{\mathbf{u}}^1 \rightarrow L_{\nu}^2$ estimate conversion). *Let $\hat{\mathbf{x}}_{\Lambda}$ be the output of any decoder and assume that \mathbf{A} has the lower robust null space property of order s . Then,*

$$\|f - \tilde{f}\|_{L_{\nu}^2} \leq \|f - f_{\Lambda}\|_{L_{\nu}^2} + \frac{2 + \rho}{\sqrt{K(s)}} \|\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda}\|_{1, \mathbf{u}} + \tau \|\mathbf{A}(\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda})\|_2.$$

where \tilde{f} is as in (13).

Proof. First, observe that

$$\|f - \tilde{f}\|_{L_{\nu}^2} = \|\mathbf{x} - \hat{\mathbf{x}}_{\Lambda}\|_2 \leq \|\mathbf{x}_{\Lambda^c}\|_2 + \|\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda}\|_2 = \|f - f_{\Lambda}\|_{L_{\nu}^2} + \|\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda}\|_2.$$

Now, choose a (not necessarily lower) set $S \subseteq \Lambda$ such that

$$\|(\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda})_{S^c}\|_2 = \inf_{\mathbf{z}: |\text{supp}(\mathbf{z})|_{\mathbf{u}} \leq K(s)} \|\mathbf{z} - (\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda})\|_2.$$

(Notice that the infimum is actually a minimum because it is defined over a finite union of linear subspaces). Employing the weighted Stechkin inequality (Lemma 5.8) and the upper bound $\|\mathbf{u}\|_{\infty}^2 \leq \frac{3}{4}K(s)$ (Lemma 5.7), we obtain

$$\|(\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda})_{S^c}\|_2 \leq \frac{\|\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda}\|_{1, \mathbf{u}}}{\sqrt{K(s) - \|\mathbf{u}_{\Lambda}\|_{\infty}^2}} \leq \frac{2\|\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda}\|_{1, \mathbf{u}}}{\sqrt{K(s)}}.$$

Moreover, the lower robust null space property implies

$$\|(\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda})_S\|_2 \leq \frac{\rho}{\sqrt{K(s)}} \|(\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda})_{S^c}\|_{1, \mathbf{u}} + \tau \|\mathbf{A}(\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda})\|_2.$$

Combining the above inequalities with the triangle inequality $\|\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda}\|_2 \leq \|(\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda})_S\|_2 + \|(\mathbf{x}_{\Lambda} - \hat{\mathbf{x}}_{\Lambda})_{S^c}\|_2$ concludes the proof. \square

5.1.2 Robustness of WQCBP

In this section, we prove that the WQCBP decoder (9) is robust under unknown error. In particular, we will assume to be in the error-blind scenario, where upper bounds of the form $\|\mathbf{e}\|_2 \leq \eta$ are not available. The results proved here generalize the robust recovery error guarantees given in [15] for the unweighted case. The recovery error estimate depends on the best s -term approximation error in lower sets, on the parameter η , on the unknown error norm $\|\mathbf{e}\|_2$, and on a tail term \mathcal{T} whose behavior is studied in Theorem 5.11. We follow the proof strategy explained in Remark 5.1.

Theorem 5.10 (Robust recovery for WQCBP). *Let $0 < \varepsilon < 1$, $0 < \delta < 1$, $2 \leq s \leq 2^{d+1}$, $\Lambda = \Lambda_{d,s}^{HC}$ be the hyperbolic cross index set defined in (16), and $\{\phi_{\mathbf{i}}\}_{\mathbf{i} \in \mathbb{N}_0^d}$ be the tensor Legendre or Chebyshev polynomial basis. Draw $\mathbf{t}_1, \dots, \mathbf{t}_m$ independently according to the corresponding measure ν , with*

$$m \gtrsim s^{\gamma} L(s, n, \delta, \varepsilon), \quad (40)$$

where γ is defined as in (20) and $L(s, n, \delta, \varepsilon)$ is the polylogarithmic factor defined in (39). Then, the following holds with probability at least $1 - \varepsilon$. For any $\eta \geq 0$ and $f \in L_{\nu}^2(D) \cap L^{\infty}(D)$ expanded as in (3), the approximation \tilde{f} defined as in (13) computed using the WQCBP decoder (9) satisfies

$$\|f - \tilde{f}\|_{L^{\infty}} \leq C_1 \sigma_{s,L}(\mathbf{x})_{1, \mathbf{u}} + s^{\gamma/2} [C_2(\eta + \|\mathbf{e}\|_2) + C_3 \mathcal{T}], \quad (41)$$

$$\|f - \tilde{f}\|_{L_{\nu}^2} \leq C_4 \frac{\sigma_{s,L}(\mathbf{x})_{1, \mathbf{u}}}{s^{\gamma/2}} + C_5(\eta + \|\mathbf{e}\|_2) + C_6 \mathcal{T} + \|f - f_{\Lambda}\|_{L_{\nu}^2}, \quad (42)$$

where f_Λ is defined as in (4),

$$\mathcal{T} = \mathcal{T}(\mathbf{A}, \Lambda, \mathbf{e}, \mathbf{u}, \eta) := \min \left\{ \frac{\|\mathbf{z}\|_{1,\mathbf{u}}}{s^{\gamma/2}} : \mathbf{z} \in \mathbb{C}^n, \|\mathbf{A}\mathbf{z} - \mathbf{e}\|_2 \leq \eta \right\}. \quad (43)$$

and with constants

$$C_1 = \frac{3+\rho}{1-\rho}, \quad C_2 = \frac{2\tau}{1-\rho}, \quad C_3 = \frac{1+\rho}{1-\rho}, \quad C_4 = \frac{2(1+\rho)(2+\rho)}{1-\rho}, \quad C_5 = \frac{\tau(9+3\rho)}{1-\rho}, \quad C_6 = C_4,$$

where $\rho = 4\delta/(1-\delta)$ and $\tau = \sqrt{1-\delta}/(1-\delta)$.

Proof. Thanks to Theorem 5.5 and Lemma 5.4, the sample complexity bound (40) guarantees that \mathbf{A} has the lower robust null space property with probability at least $1-\varepsilon$. As a consequence, Theorem 5.6 holds with probability at least $1-\varepsilon$.

Now, we observe that, on the one hand,

$$\|\mathbf{A}(\mathbf{x}_\Lambda - \hat{\mathbf{x}}_\Lambda)\|_2 \leq \|\mathbf{A}\mathbf{x}_\Lambda - \mathbf{y}\|_2 + \|\mathbf{A}\hat{\mathbf{x}}_\Lambda - \mathbf{y}\|_2 \leq \|\mathbf{e}\|_2 + \eta \quad (44)$$

and, on the other hand,

$$\begin{aligned} \|\hat{\mathbf{x}}_\Lambda\|_{1,\mathbf{u}} - \|\mathbf{x}_\Lambda\|_{1,\mathbf{u}} &= \min\{\|\mathbf{z}\|_{1,\mathbf{u}} : \mathbf{z} \in \mathbb{C}^n, \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2 \leq \eta\} - \|\mathbf{x}_\Lambda\|_{1,\mathbf{u}} \\ &\leq \min\{\|\mathbf{z} - \mathbf{x}_\Lambda\|_{1,\mathbf{u}} : \mathbf{z} \in \mathbb{C}^n, \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2 \leq \eta\} \\ &= \min\{\|\mathbf{z}\|_{1,\mathbf{u}} : \mathbf{z} \in \mathbb{C}^n, \|\mathbf{A}\mathbf{z} - \mathbf{e}\|_2 \leq \eta\} \\ &= s^{\gamma/2} \mathcal{T}(\mathbf{A}, \Lambda, \mathbf{e}, \mathbf{u}, \eta). \end{aligned}$$

Noting that $\sigma_{s,L}(\mathbf{x}_\Lambda)_{1,\mathbf{u}} = \sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}$ (since Λ is the union of all lower sets of cardinality s), that $K(s) \leq s^\gamma$ (due to Lemma 2.2), and employing Theorem 5.6 with $\mathbf{z} = \mathbf{x}_\Lambda$ and $\hat{\mathbf{z}} = \hat{\mathbf{x}}_\Lambda$, we obtain

$$\|\mathbf{x}_\Lambda - \hat{\mathbf{x}}_\Lambda\|_{1,\mathbf{u}} \leq \frac{2(1+\rho)}{1-\rho} \sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}} + s^{\gamma/2} \left[\frac{2\tau}{1-\rho} (\|\mathbf{e}\|_2 + \eta) + \frac{1+\rho}{1-\rho} \mathcal{T} \right]. \quad (45)$$

Using $\|\mathbf{x} - \mathbf{x}_\Lambda\|_{1,\mathbf{u}} \leq \sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}$ and plugging (45) into

$$\|f - \tilde{f}\|_{L^\infty} \leq \|\mathbf{x} - \hat{\mathbf{x}}_\Lambda\|_{1,\mathbf{u}} \leq \|\mathbf{x} - \mathbf{x}_\Lambda\|_{1,\mathbf{u}} + \|\mathbf{x}_\Lambda - \hat{\mathbf{x}}_\Lambda\|_{1,\mathbf{u}},$$

gives (41).

Finally, we employ the $\ell_{\mathbf{u}}^1 \rightarrow L_\nu^2$ estimate conversion result (Lemma 5.9) combined with (44), (45), and with the fact that $K(s) \geq \frac{1}{4}s^\gamma$ (Lemma 2.2), to obtain (42). \square

In the following result, we give an upper bound to the tail term \mathcal{T} defined in (43) in terms of the m^{th} singular value (in decreasing order) of the matrix $\sqrt{\frac{m}{n}}\mathbf{A}^*$, denoted as $\sigma_m(\sqrt{\frac{m}{n}}\mathbf{A}^*)$.

Theorem 5.11 (Tail term bound). *Consider the setup of Theorem 5.10 with*

$$m \asymp s^\gamma L(s, n, \delta, \varepsilon), \quad (46)$$

and let $\mathcal{T} = \mathcal{T}(\mathbf{A}, \Lambda, \mathbf{e}, \mathbf{u}, \eta)$ be as in (43). Then,

$$\mathcal{T} \lesssim \frac{s^{\alpha/2} \sqrt{L}}{\sigma_m(\sqrt{\frac{m}{n}}\mathbf{A}^*)} \max\{\|\mathbf{e}\|_2 - \eta, 0\}, \quad (47)$$

where L is as in (39) and $\alpha = 1, 2$ in the Chebyshev or Legendre case, respectively.

Proof. Notice that if the rank of \mathbf{A} is not full, then $\sigma_m(\sqrt{\frac{m}{n}}\mathbf{A}^*) = 0$ and (47) is trivially satisfied. Hence, we assume that \mathbf{A} has full rank. Also if $\eta \geq \|e\|_2$ the result holds trivially. Then, we also suppose that $\eta < \|e\|_2$.

Since $\|e\|_2 \neq 0$, we can define the ansatz $\mathbf{z} := (1 - \eta/\|e\|_2)\mathbf{A}^\dagger e$, where $\mathbf{A}^\dagger = \mathbf{A}^*(\mathbf{A}\mathbf{A}^*)^{-1}$ is the pseudoinverse. Then \mathbf{z} satisfies $\|\mathbf{A}\mathbf{z} - e\|_2 = \eta$, and hence, recalling the definition (43) of \mathcal{T} , we have

$$s^{\gamma/2}\mathcal{T} \leq \|\mathbf{z}\|_{1,\mathbf{u}} \leq \sqrt{|\Lambda|_{\mathbf{u}}}\|\mathbf{z}\|_2 \leq \frac{\sqrt{|\Lambda|_{\mathbf{u}}}}{\sigma_m(\mathbf{A}^*)}(\|e\|_2 - \eta).$$

Equation (46) implies that $\sqrt{\frac{m}{s^\gamma}} \lesssim \sqrt{L}$, and therefore

$$\mathcal{T} \lesssim \sqrt{\frac{|\Lambda|_{\mathbf{u}}}{n}} \frac{\sqrt{L}}{\sigma_m(\sqrt{\frac{m}{n}}\mathbf{A}^*)} (\|e\|_2 - \eta). \quad (48)$$

The last step is to estimate $|\Lambda|_{\mathbf{u}}$ using the explicit formulae (15) for the intrinsic weights. For the Chebyshev case, we have

$$|\Lambda|_{\mathbf{u}} = \sum_{\mathbf{i} \in \Lambda} 2^{\|\mathbf{i}\|_0} \leq \sum_{\mathbf{i} \in \Lambda} \prod_{\ell=1}^d (i_\ell + 1) \leq s \sum_{\mathbf{i} \in \Lambda} 1 = sn$$

where in the penultimate step we used the definition of the hyperbolic cross (16). For the Legendre case, we have

$$|\Lambda|_{\mathbf{u}} = \sum_{\mathbf{i} \in \Lambda} \prod_{\ell=1}^d (2i_\ell + 1) \leq \sum_{\mathbf{i} \in \Lambda} 2^{\|\mathbf{i}\|_0} \prod_{\ell=1}^d (i_\ell + 1) \leq s^2 n.$$

This completes the proof. \square

Theorems 5.10 & 5.11 show that the robustness of WQCBP is implied by inequalities of the form $\sigma_m(\sqrt{\frac{m}{n}}\mathbf{A}^*) \gtrsim 1$. When $d = 1$, this can be achieved by resorting to the spectral theory of random matrices with heavy-tailed rows [15]. Showing this type of inequality when $d > 1$ is still an open problem. However, it is not difficult to show that (see [4, Lemma 3])

$$\lambda_{\min}\left(\mathbb{E}\left[\frac{m}{n}\mathbf{A}\mathbf{A}^*\right]\right) = 1 - \frac{1}{n}.$$

By inspecting the proof of Theorem 5.11, we notice that

$$\mathcal{T} \lesssim \mathcal{Q}\sqrt{L} \max\{\|e\|_2 - \eta, 0\},$$

where

$$\mathcal{Q} = \mathcal{Q}(\mathbf{A}, \Lambda, \mathbf{u}) := \sqrt{\frac{|\Lambda|_{\mathbf{u}}}{n}} \frac{1}{\sigma_m(\sqrt{\frac{m}{n}}\mathbf{A}^*)}. \quad (49)$$

This constant can be easily estimated numerically and turns out to have moderate size (see [4]).

To conclude, we notice that Theorem 5.10 implies an analogous result in the error-aware scenario, i.e. when $\|e\|_2 \leq \eta$.

Theorem 5.12 (Robust recovery of WQCBP in the error-aware setting). *Consider the same setup of Theorem 5.10 and let $\|e\|_2 \leq \eta$. Then, with probability at least $1 - \varepsilon$ the following inequalities hold:*

$$\begin{aligned} \|f - \tilde{f}\|_{L^\infty} &\lesssim \sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}} + s^{\gamma/2}\eta, \\ \|f - \tilde{f}\|_{L_v^2} &\lesssim \frac{\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}}{s^{\gamma/2}} + \eta + \|f - f_\Lambda\|_{L_v^2}. \end{aligned}$$

5.1.3 Robustness of Wlasso

In this section, we prove that the Wlasso decoder is robust under unknown error, when the tuning parameter is chosen proportionally to the ratio $\sqrt{K(s)}/\|e\|_2$.

Theorem 5.13 (Wlasso robust recovery). *Let $0 < \varepsilon < 1$, $0 < \delta < 1$, $2 \leq s \leq 2^{d+1}$, $\Lambda = \Lambda_{d,s}^{HC}$ be the hyperbolic cross index set defined in (16), and $\{\phi_i\}_{i \in \mathbb{N}_0^d}$ be the tensor Legendre or Chebyshev polynomial basis. Draw $\mathbf{t}_1, \dots, \mathbf{t}_m$ independently according to the corresponding measure ν , with*

$$m \gtrsim s^\gamma L(s, n, \delta, \varepsilon),$$

where γ is defined as in (20) and $L(s, n, \delta, \varepsilon)$ is the polylogarithmic factor defined in (39). Then, the following holds with probability at least $1 - \varepsilon$. For any $f \in L_\nu^2(D) \cap L^\infty(D)$ expanded as in (3), the approximation \tilde{f} defined as in (13) computed using the Wlasso decoder (10) with tuning parameter

$$\lambda = \theta \frac{\sqrt{K(s)}}{\|e\|_2}, \quad \text{with } \theta > 0, \quad (50)$$

satisfies

$$\|f - \tilde{f}\|_{L^\infty} \leq C_1 \sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}} + C_2 s^{\gamma/2} \|e\|_2, \quad (51)$$

$$\|f - \tilde{f}\|_{L_\nu^2} \leq C_3 \frac{\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}}{s^{\gamma/2}} + C_4 \|e\|_2 + \|f - f_\Lambda\|_{L_\nu^2}, \quad (52)$$

where f_Λ is defined as in (4), and the constants are

$$C_1 = \frac{3 + \rho}{1 - \rho}, \quad C_2 = \frac{1}{1 - \rho} \left[\frac{\tau^2}{(1 + \rho)\theta} + (1 + \rho)\theta + 2\tau \right],$$

and

$$C_3 = \frac{4(1 + \rho)(2 + \rho)}{1 - \rho} \quad C_4 = \frac{1}{1 - \rho} \left[\frac{(5 + \rho)^2 \tau^2}{4(1 + \rho)(2 + \rho)\theta} + (1 + \rho)(2 + \rho)\theta + (5 + \rho)\tau \right],$$

where $\rho = 4\delta/(1 - \delta)$, $\tau = \sqrt{1 - \delta}/(1 - \delta)$.

Proof. First, notice that (10) can be reformulated in augmented form as

$$(\hat{\mathbf{x}}_\Lambda, \hat{\mathbf{e}}) := \arg \min_{(\mathbf{z}, \mathbf{d}) \in \mathbb{C}^n \times \mathbb{C}^m} \|\mathbf{z}\|_{1,\mathbf{u}} + \lambda \|\mathbf{d}\|_2^2 \text{ s.t. } \mathbf{A}\mathbf{z} + \mathbf{d} = \mathbf{y}.$$

Moreover, Theorem 5.6 (that holds with probability at least $1 - \varepsilon$) implies

$$\|\mathbf{x}_\Lambda - \hat{\mathbf{x}}_\Lambda\|_{1,\mathbf{u}} \leq \frac{2(1 + \rho)}{1 - \rho} \sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}} + \frac{\xi}{1 - \rho}, \quad (53)$$

where

$$\xi := (1 + \rho)(\|\hat{\mathbf{x}}_\Lambda\|_{1,\mathbf{u}} - \|\mathbf{x}_\Lambda\|_{1,\mathbf{u}}) + 2\tau \sqrt{K(s)} \|\mathbf{A}(\hat{\mathbf{x}}_\Lambda - \mathbf{x}_\Lambda)\|_2. \quad (54)$$

Since $\|\mathbf{A}(\hat{\mathbf{x}}_\Lambda - \mathbf{x}_\Lambda)\|_2 \leq \|\hat{\mathbf{e}}\|_2 + \|e\|_2$, we estimate

$$\xi \leq (1 + \rho)\|\hat{\mathbf{x}}_\Lambda\|_{1,\mathbf{u}} + 2\tau \sqrt{K(s)} \|\hat{\mathbf{e}}\|_2 - (1 + \rho)\|\mathbf{x}_\Lambda\|_{1,\mathbf{u}} + 2\tau \sqrt{K(s)} \|e\|_2.$$

Now, we employ Young's inequality $ab \leq \omega a^2 + \frac{b^2}{4\omega}$, with $\omega = (1 + \rho)\lambda$, $a = \|\hat{\mathbf{e}}\|_2$, and $b = 2\tau \sqrt{K(s)}$. This, combined with the optimality of $(\hat{\mathbf{x}}_\Lambda, \hat{\mathbf{e}})$ yields

$$\xi \leq \frac{\tau^2 K(s)}{(1 + \rho)\lambda} + (1 + \rho)\lambda \|e\|_2^2 + 2\tau \sqrt{K(s)} \|e\|_2.$$

Now, plugging (50) into the above relation, we see that

$$\xi \leq \left[\frac{\tau^2}{(1+\rho)\theta} + (1+\rho)\theta + 2\tau \right] \sqrt{K(s)} \|\mathbf{e}\|_2.$$

Finally, recalling Lemma 2.2, from (53) we obtain (51) as in Theorem 5.10.

Now, Lemma 5.9 combined with (53) yields

$$\|f - \tilde{f}\|_{L_\nu^2} \leq \|f - f_\Lambda\|_{L_\nu^2} + \frac{2(1+\rho)(2+\rho)}{1-\rho} \frac{\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}}{\sqrt{K(s)}} + \frac{\zeta}{(1-\rho)\sqrt{K(s)}} \quad (55)$$

where

$$\zeta := (1+\rho)(2+\rho)(\|\hat{\mathbf{x}}\|_{1,\mathbf{u}} - \|\mathbf{x}\|_{1,\mathbf{u}}) + (5+\rho)\tau\sqrt{K(s)}\|\mathbf{A}(\hat{\mathbf{x}}_\Lambda - \mathbf{x}_\Lambda)\|_2 \quad (56)$$

Employing Young's inequality with $\omega = (1+\rho)(2+\rho)\lambda$, $a = \|\hat{\mathbf{e}}\|_2$, and $(5+\rho)\tau\sqrt{K(s)}$ and using analogous manipulations as before, we have

$$\zeta \leq \left[\frac{(5+\rho)^2\tau^2}{4(1+\rho)(2+\rho)\theta} + (1+\rho)(2+\rho)\theta + (5+\rho)\tau \right] \sqrt{K(s)} \|\mathbf{e}\|_2.$$

Combining the above inequality with (55) and recalling Lemma 2.2 gives (52), as required. \square

5.1.4 Robustness of WSR-LASSO

We prove a robustness under unknown error for the WSR-LASSO decoder (11). Our theory suggests that the tuning parameter should be chosen directly proportional to the quantity $\sqrt{K(s)}$. This choice is advantageous, since it is independent of the unknown error \mathbf{e} .

Theorem 5.14 (WSR-LASSO robust recovery). *Let $0 < \varepsilon < 1$, $0 < \delta < 1$, $2 \leq s \leq 2^{d+1}$, $\Lambda = \Lambda_{d,s}^{HC}$ be the hyperbolic cross index set defined in (16), and $\{\phi_{\mathbf{i}}\}_{\mathbf{i} \in \mathbb{N}_0^d}$ be the tensor Legendre or Chebyshev polynomial basis. Draw $\mathbf{t}_1, \dots, \mathbf{t}_m$ independently according to the corresponding measure ν , with*

$$m \gtrsim s^\gamma L(s, n, \delta, \varepsilon),$$

where γ is defined as in (20) and $L(s, n, \delta, \varepsilon)$ is the polylogarithmic factor defined in (39). Then, the following holds with probability at least $1 - \varepsilon$. For any $f \in L_\nu^2(D) \cap L^\infty(D)$ expanded as in (3), the approximation \tilde{f} defined as in (13) computed using the WSR-LASSO decoder (11) with tuning parameter

$$\lambda = \theta\sqrt{K(s)}, \quad \text{with } \theta \geq \frac{(5+\rho)\tau}{(1+\rho)(2+\rho)} \quad (57)$$

where $\rho = 4\delta/(1-\delta)$ and $\tau = \sqrt{1-\delta}/(1-\delta)$, satisfies

$$\|f - \tilde{f}\|_{L^\infty} \leq C_1 \sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}} + C_2 s^{\gamma/2} \|\mathbf{e}\|_2, \quad (58)$$

$$\|f - \tilde{f}\|_{L_\nu^2} \leq C_3 \frac{\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}}{s^{\gamma/2}} + C_4 \|\mathbf{e}\|_2 + \|f - f_\Lambda\|_{L_\nu^2}, \quad (59)$$

where f_Λ is defined as in (4) and the constants are

$$C_1 = \frac{3+\rho}{1-\rho}, \quad C_2 = \frac{(1+\rho)\theta + 2\tau}{1-\rho}, \quad C_3 = \frac{4(1+\rho)(2+\rho)}{1-\rho}, \quad C_4 = \frac{(1+\rho)(2+\rho)\theta + (5+\rho)\tau}{1-\rho}.$$

Proof. The proof is similar to that of Theorem 5.13. First, observe that (11) admits the equivalent augmented formulation

$$(\hat{\mathbf{x}}_\Lambda, \hat{\mathbf{e}}) = \arg \min_{(\mathbf{z}, \mathbf{d}) \in \mathbb{C}^n \times \mathbb{C}^m} \|\mathbf{z}\|_{1,\mathbf{u}} + \lambda \|\mathbf{d}\|_2 \text{ s.t. } \mathbf{A}\mathbf{z} + \mathbf{d} = \mathbf{y}.$$

Moreover, Theorem 5.6 (that holds with probability at least $1 - \varepsilon$) implies the ℓ_u^1 estimate (53). Now, we observe that $\rho < 1$ and (57) yield

$$\lambda \geq \frac{(5 + \rho)\tau\sqrt{K(s)}}{(1 + \rho)(2 + \rho)} \geq \frac{2\tau\sqrt{K(s)}}{1 + \rho}.$$

Therefore, using the above inequality and the optimality of $(\hat{\mathbf{x}}_\Lambda, \hat{\mathbf{e}})$, we obtain

$$\begin{aligned} \xi &\leq (1 + \rho)(\|\hat{\mathbf{x}}_\Lambda\|_{1,\mathbf{u}} + \lambda\|\hat{\mathbf{e}}\|_2) + 2\tau\sqrt{K(s)}\|\mathbf{e}\|_2 - (1 + \rho)\|\mathbf{x}_\Lambda\|_{1,\mathbf{u}} \\ &\leq [(1 + \rho)\lambda + 2\tau\sqrt{K(s)}]\|\mathbf{e}\|_2 = [(1 + \rho)\theta + 2\tau]\sqrt{K(s)}\|\mathbf{e}\|_2, \end{aligned}$$

where ξ is defined as in (54). Plugging the above inequality into (53) gives (58).

Again, analogously to the proof of Theorem 5.13, we have that the L_ν^2 estimate (55) holds. Then, using (57) and an argument analogous to that employed to bound ξ , we obtain

$$\zeta \leq [(1 + \rho)(2 + \rho)\theta + (5 + \rho)\tau]\sqrt{K(s)}\|\mathbf{e}\|_2,$$

where ζ is defined as in (56). Recalling Lemma 2.2, we have (59). The proof is thus complete. \square

5.2 WLAD-LASSO

The proof strategy to show the robustness of WLAD-LASSO under unknown error still follows the guidelines of Remark 5.1. Yet, in this case we need to introduce a further level of technical difficulty, dealing with the framework of sparsity in levels, introduced in [5]. In particular, we will make use of a weighted version of the 2-level sparsity. We introduce this notion and some related technical lemmas in §5.2.1. Then, we show the robustness of WLAD-LASSO in §5.2.2.

For the sake of generality, the results in this section are proved for a variant of (12) with weighted ℓ^1 penalization of the data-fidelity term

$$\hat{\mathbf{x}}_\Lambda := \arg \min_{\mathbf{z} \in \mathbb{C}^n} \|\mathbf{z}\|_{1,\mathbf{u}} + \lambda \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_{1,\mathbf{v}}, \quad (60)$$

where $\mathbf{v} \in \mathbb{R}^m$ is such that $v_i \geq 1$ for every $i \in [m]$.

5.2.1 Preliminaries

We first introduce some notation regarding the weighted 2-level sparsity framework [5].

Definition 5.15 (Weighted 2-level sparsity notation). Let $n, p \in \mathbb{N}$. Consider a vector $\mathbf{z} \in \mathbb{C}^{n+p}$ and weights $\mathbf{w} \in \mathbb{R}^{n+p}$, partitioned as

$$\mathbf{z} = \begin{bmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \end{bmatrix}, \quad \mathbf{w} = \begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \end{bmatrix},$$

such that $\mathbf{z}_1 \in \mathbb{C}^n$, $\mathbf{z}_2 \in \mathbb{C}^p$, $\mathbf{w}_1 \in \mathbb{R}^n$, and $\mathbf{w}_2 \in \mathbb{R}^p$. Then, given $\mathbf{J} := (J_1, J_2) \in \mathbb{R}^2$, \mathbf{z} is said to be *weighted \mathbf{J} -sparse* if

$$|\text{supp}(\mathbf{z}_i)|_{\mathbf{w}_i} \leq J_i, \quad \forall i = 1, 2.$$

The *weighted best \mathbf{J} -term approximation error* to \mathbf{z} is

$$\sigma_{\mathbf{J}}(\mathbf{z})_{1,\mathbf{w}} := \inf_{\mathbf{z}': |\text{supp}(\mathbf{z}')|_{\mathbf{w}_i} \leq J_i} \|\mathbf{z} - \mathbf{z}'\|_{1,\mathbf{w}}.$$

We also introduce a *scale* $\lambda > 0$ that will allow us to switch from the two- to the one-level formalism.

Definition 5.16 (Scaled weights and scaled sparsity). Let $\lambda > 0$ be a *scale*, then for every $\mathbf{w} = \begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \end{bmatrix} \in \mathbb{C}^{n+p}$ and $\mathbf{J} = (J_1, J_2) \in \mathbb{R}^2$, we define the *scaled weights* \mathbf{w}_λ and the *scaled sparsity* J_λ as

$$\mathbf{w}_\lambda := \begin{bmatrix} \mathbf{w}_1 \\ \lambda \mathbf{w}_2 \end{bmatrix}, \quad J_\lambda = J_1 + \lambda^2 J_2. \quad (61)$$

Notice that if \mathbf{z} is \mathbf{J} -sparse, then

$$\|\mathbf{z}\|_{1, \mathbf{w}_\lambda} \leq \sqrt{J_\lambda} \|\mathbf{z}\|_2. \quad (62)$$

Using the notation above, the WLAD-LASSO optimization program (60) admits an equivalent augmented formulation as a weighted basis pursuit program, i.e.

$$\begin{bmatrix} \hat{\mathbf{x}}_\Lambda \\ \hat{\mathbf{e}} \end{bmatrix} = \arg \min_{\mathbf{z} \in \mathbb{C}^{n+m}} \|\mathbf{z}\|_{1, \mathbf{w}_\lambda} \text{ s.t. } \mathbf{M}\mathbf{z} = \mathbf{y}, \quad (63)$$

where $\mathbf{w} := \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix}$, $\mathbf{M} := [\mathbf{A} \quad \mathbf{I}]$ and \mathbf{w}_λ is defined as in (61).

We are now in a position to define the 2-level weighted robust null space property.

Definition 5.17 (2-level weighted robust null space property). A matrix $\mathbf{M} \in \mathbb{C}^{m \times (N+p)}$ is said to satisfy the *2-level weighted robust null space property* of scale $\lambda > 0$ and order $\mathbf{J} = (J_1, J_2)$ if

$$\|\mathbf{z}_{S_1 \cup S_2}\|_2 \leq \frac{\rho}{\sqrt{J_\lambda}} \|\mathbf{z}_{(S_1 \cup S_2)^c}\|_{1, \mathbf{w}_\lambda} + \tau \|\mathbf{M}\mathbf{z}\|_2,$$

for every $S_1 \subseteq [n], S_2 \subseteq n + [p]$, such that $|S_i|_{\mathbf{w}_i} \leq J_i$, for $i = 1, 2$.

Notice that in the one-level case (i.e., when $p = 0$) and with $J_1 = K(s)$ Definition 5.17 and Definition 5.3 are equivalent.

We now prove a result analogous to Theorem 5.6. It is an upper bound to the $\ell_{\mathbf{w}_\lambda}^1$ distance between two arbitrary finite-dimensional vectors under the weighted robust null space property. This result is the 2-level generalization of [51, Theorem 4.2] and is based on the same proof strategy.

Theorem 5.18 ($\ell_{\mathbf{w}_\lambda}^1$ distance bound). *If the matrix $\mathbf{M} \in \mathbb{C}^{m \times (N+p)}$ has the 2-level weighted robust null space property, then for every $\mathbf{z}, \hat{\mathbf{z}} \in \mathbb{C}^{n+p}$ the following holds:*

$$\|\mathbf{z} - \hat{\mathbf{z}}\|_{1, \mathbf{w}_\lambda} \leq \frac{1 + \rho}{1 - \rho} (2\sigma_{\mathbf{J}}(\mathbf{z})_{1, \mathbf{w}_\lambda} + \|\hat{\mathbf{z}}\|_{1, \mathbf{w}_\lambda} - \|\mathbf{z}\|_{1, \mathbf{w}_\lambda}) + \frac{2\tau\sqrt{J_\lambda}}{1 - \rho} \|\mathbf{M}(\mathbf{z} - \hat{\mathbf{z}})\|_2. \quad (64)$$

Proof. Choose $S_1 \subseteq [n]$ and $S_2 \subseteq [p]$ such that $\sigma_{\mathbf{J}}(\mathbf{z})_{1, \mathbf{w}_\lambda} = \|\mathbf{z} - \mathbf{z}_T\|_{1, \mathbf{w}_\lambda} = \|\mathbf{z}_{T^c}\|_{1, \mathbf{w}_\lambda}$ and define $T := S_1 \cup S_2$. Using the 2-level weighted robust null space property and recalling (62), we obtain

$$\|(\mathbf{z} - \hat{\mathbf{z}})_T\|_{1, \mathbf{w}_\lambda} \leq \rho \|(\mathbf{z} - \hat{\mathbf{z}})_{T^c}\|_{1, \mathbf{w}_\lambda} + \tau\sqrt{J_\lambda} \|\mathbf{M}(\mathbf{z} - \hat{\mathbf{z}})\|_2. \quad (65)$$

Then, we estimate

$$\begin{aligned} \|(\mathbf{z} - \hat{\mathbf{z}})_{T^c}\|_{1, \mathbf{w}_\lambda} &\leq \|\mathbf{z}_{T^c}\|_{1, \mathbf{w}_\lambda} + \|\hat{\mathbf{z}}_{T^c}\|_{1, \mathbf{w}_\lambda} \\ &= 2\|\mathbf{z}_{T^c}\|_{1, \mathbf{w}_\lambda} + \|\mathbf{z}_T\|_{1, \mathbf{w}_\lambda} - \|\hat{\mathbf{z}}_T\|_{1, \mathbf{w}_\lambda} + \|\hat{\mathbf{z}}\|_{1, \mathbf{w}_\lambda} - \|\mathbf{z}\|_{1, \mathbf{w}_\lambda} \\ &\leq 2\sigma_{\mathbf{J}}(\mathbf{z})_{1, \mathbf{w}_\lambda} + \|(\mathbf{z} - \hat{\mathbf{z}})_T\|_{1, \mathbf{w}_\lambda} + \|\hat{\mathbf{z}}\|_{1, \mathbf{w}_\lambda} - \|\mathbf{z}\|_{1, \mathbf{w}_\lambda}. \end{aligned}$$

Plugging (65) into the inequality above and solving for $\|(\mathbf{z} - \hat{\mathbf{z}})_{T^c}\|_{1, \mathbf{w}_\lambda}$, we see that

$$\|(\mathbf{z} - \hat{\mathbf{z}})_{T^c}\|_{1, \mathbf{w}_\lambda} \leq \frac{2\sigma_{\mathbf{J}}(\mathbf{z})_{1, \mathbf{w}_\lambda} + \tau\sqrt{J_\lambda} \|\mathbf{M}(\mathbf{z} - \hat{\mathbf{z}})\|_2 + \|\hat{\mathbf{z}}\|_{1, \mathbf{w}_\lambda} - \|\mathbf{z}\|_{1, \mathbf{w}_\lambda}}{1 - \rho}. \quad (66)$$

Combining (65) and (66) with the triangle inequality $\|\mathbf{z} - \hat{\mathbf{z}}\|_{1, \mathbf{w}_\lambda} \leq \|(\mathbf{z} - \hat{\mathbf{z}})_T\|_{1, \mathbf{w}_\lambda} + \|(\mathbf{z} - \hat{\mathbf{z}})_{T^c}\|_{1, \mathbf{w}_\lambda}$ concludes the proof. \square

Before proving Theorem 5.25, we need to introduce some more technical elements to the picture. First, we consider the 2-level weighted restricted isometry property (a weighted version of the restricted isometry property in levels, introduced in [11], also studied in [3] for sparse corruptions).

Definition 5.19 (2-level weighted restricted isometry property). Let $\mathbf{J} \in \mathbb{R}^2$, and $\mathbf{w} \in \mathbb{R}^{n+p}$ partitioned as in Definition 5.15. Then, a matrix $\mathbf{M} \in \mathbb{C}^{m \times (n+p)}$ is said to have the *2-level weighted restricted isometry property* of order \mathbf{J} if there exists a constant $0 < \delta < 1$ such that

$$(1 - \delta)\|\mathbf{z}\|_2^2 \leq \|\mathbf{M}\mathbf{z}\|_2^2 \leq (1 + \delta)\|\mathbf{z}\|_2^2, \quad \forall \mathbf{z} \in \mathbb{C}^{n+p} : |\text{supp}(\mathbf{z}_i)|_{\mathbf{w}_i} \leq J_i, \forall i = 1, 2.$$

The smallest constant such that this property holds is called the \mathbf{J}^{th} *2-level weighted restricted isometry constant* of \mathbf{M} and it is denoted as $\delta_{\mathbf{J}}$.

As a consequence of the definition above, we have the following technical result.

Lemma 5.20. Consider $\mathbf{z}, \hat{\mathbf{z}} \in \mathbb{C}^{n+p}$, weights $\mathbf{w} \in \mathbb{R}^{n+p}$, and $\mathbf{J}, \hat{\mathbf{J}} \in \mathbb{R}^2$ partitioned as in Definition 5.15 and such that $|\text{supp}(\mathbf{z}_i)|_{\mathbf{w}_i} \leq J_i$ and $|\text{supp}(\hat{\mathbf{z}}_i)|_{\mathbf{w}_i} \leq \hat{J}_i$, for $i = 1, 2$. Moreover, assume that $\text{supp}(\mathbf{z}_1) \cap \text{supp}(\hat{\mathbf{z}}_1) = \text{supp}(\mathbf{z}_2) \cap \text{supp}(\hat{\mathbf{z}}_2) = \emptyset$. Then, it holds

$$|\langle \mathbf{M}\mathbf{z}, \mathbf{M}\hat{\mathbf{z}} \rangle| \leq \delta_{\mathbf{J}+\hat{\mathbf{J}}} \|\mathbf{z}\|_2 \|\hat{\mathbf{z}}\|_2, \quad \text{with } \mathbf{M} = [\mathbf{A} \quad \mathbf{I}].$$

Proof. Let $S_1 := \text{supp}(\mathbf{z}_1) \cup \text{supp}(\hat{\mathbf{z}}_1)$ and $S_2 := \text{supp}(\mathbf{z}_2) \cup \text{supp}(\hat{\mathbf{z}}_2)$. Then $|S_i|_{\mathbf{w}_i} \leq J_i + \hat{J}_i$, for $i = 1, 2$. Using the disjointedness of the supports of \mathbf{z}_1 and $\hat{\mathbf{z}}_1$ and of those of \mathbf{z}_2 and $\hat{\mathbf{z}}_2$, we have

$$\begin{aligned} |\langle \mathbf{M}\mathbf{z}, \mathbf{M}\hat{\mathbf{z}} \rangle| &= |\langle \mathbf{M}\mathbf{z}, \mathbf{M}\hat{\mathbf{z}} \rangle - \langle \mathbf{z}, \hat{\mathbf{z}} \rangle| \\ &= |\langle [\mathbf{A}_{S_1} \quad \mathbf{I}_{S_2}] \mathbf{z}_{S_1 \cup S_2} [\mathbf{A}_{S_1} \quad \mathbf{I}_{S_2}] \hat{\mathbf{z}}_{S_1 \cup S_2} \rangle - \langle \mathbf{z}_{S_1 \cup S_2}, \hat{\mathbf{z}}_{S_1 \cup S_2} \rangle| \\ &= \left| \left\langle \left([\mathbf{A}_{S_1} \quad \mathbf{I}_{S_2}]^* [\mathbf{A}_{S_1} \quad \mathbf{I}_{S_2}] - \mathbf{I} \right) \mathbf{z}_{S_1 \cup S_2}, \hat{\mathbf{z}}_{S_1 \cup S_2} \right\rangle \right| \\ &\leq \| [\mathbf{A}_{S_1} \quad \mathbf{I}_{S_2}]^* [\mathbf{A}_{S_1} \quad \mathbf{I}_{S_2}] - \mathbf{I} \|_{2 \rightarrow 2} \|\mathbf{z}_{S_1 \cup S_2}\|_2 \|\hat{\mathbf{z}}_{S_1 \cup S_2}\|_2 \\ &\leq \delta_{\mathbf{J}+\hat{\mathbf{J}}} \|\mathbf{z}\|_2 \|\hat{\mathbf{z}}\|_2. \end{aligned}$$

This completes the proof. \square

In view of Remark 5.1, we now prove that the restricted isometry property implies the robust null space property in the 2-level weighted setting.

Theorem 5.21 (2-level weighted restricted isometry property \Rightarrow 2-level weighted robust null space property). Given $\mathbf{J} = (K, H) \in \mathbb{R}^2$, any matrix $\mathbf{M} \in \mathbb{C}^{m \times (n+p)}$ of the form $\mathbf{M} = [\mathbf{A} \quad \mathbf{I}]$ with $(3\mathbf{J})^{th}$ weighted restricted isometry constant

$$\delta_{3\mathbf{J}} < \frac{1}{1 + 4\Theta}, \quad \text{with } \Theta := \frac{\sqrt{K + \lambda^2 H}}{\min\{\sqrt{K}, \lambda\sqrt{H}\}}, \quad (67)$$

with respect to weights $\mathbf{w} = \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} \in \mathbb{R}^{n+p}$ such that $\mathbf{J} \geq \frac{4}{3}(\|\mathbf{u}\|_\infty^2, \|\mathbf{v}\|_\infty^2)$ satisfies the 2-level weighted robust null space property of order \mathbf{J} with constants

$$\rho = \frac{4\delta_{3\mathbf{J}}\Theta}{(1 - \delta_{3\mathbf{J}})}, \quad \tau = \frac{\sqrt{1 + \delta_{3\mathbf{J}}}}{1 - \delta_{3\mathbf{J}}}. \quad (68)$$

Proof. Consider $\mathbf{x} \in \mathbb{C}^n$ and $S \subset [n]$ with $|S|_{\mathbf{u}} \leq K$. We partition $S^c = [n] \setminus S$ into disjoint blocks S_1, S_2, \dots with $K - \|\mathbf{u}\|_\infty^2 \leq |S_\ell|_{\mathbf{u}} \leq K$ according to the non-increasing rearrangement of $(|x_j|u_j^{-1})_{j \in S^c}$. As a result, we have $|x_j|u_j^{-1} \leq |x_k|u_k^{-1}$ for all $j \in S_\ell, k \in S_{\ell-1}$, and $\ell \geq 2$. Similarly, we let $\mathbf{e} \in \mathbb{C}^p$ and $T \subset [p]$ with $|T|_{\mathbf{v}_2} \leq H$. We partition T^c into blocks T_1, T_2, \dots with $H - \|\mathbf{v}\|_\infty^2 \leq |S_\ell|_{\mathbf{v}} \leq H$ according to the non-increasing rearrangement of $(|e_j|v_j^{-1})_{j \in T^c}$. We have $|e_j|v_j^{-1} \leq |e_k|v_k^{-1}$ for all $j \in T_\ell, k \in T_{\ell-1}$,

and $\ell \geq 2$. Using and the 2-level weighted restricted isometry property, we have

$$\begin{aligned}
\left\| \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix} \right\|_2^2 &\leq \frac{1}{1 - \delta_{2J}} \left\| M \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix} \right\|_2^2 \\
&= \frac{1}{1 - \delta_{2J}} \left\langle M \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix}, M \begin{bmatrix} \mathbf{x} \\ \mathbf{e} \end{bmatrix} - \sum_{\ell \geq 2} M \begin{bmatrix} \mathbf{x}_{S_\ell} \\ \mathbf{e}_{T_\ell} \end{bmatrix} \right\rangle \\
&= \frac{1}{1 - \delta_{2J}} \left(\left\langle M \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix}, M \begin{bmatrix} \mathbf{x} \\ \mathbf{e} \end{bmatrix} \right\rangle - \sum_{\ell \geq 2} \left\langle M \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix}, M \begin{bmatrix} \mathbf{x}_{S_\ell} \\ \mathbf{e}_{T_\ell} \end{bmatrix} \right\rangle \right)
\end{aligned}$$

Then, employing again the restricted isometry property, Lemma 5.20, and the Cauchy-Schwarz inequality, we see that

$$\begin{aligned}
\left\| \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix} \right\|_2^2 &\leq \frac{1}{1 - \delta_{2J}} \left(\left\| M \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix} \right\|_2 \|\mathbf{Ax} + \mathbf{e}\|_2 + \delta_{3J} \left\| \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix} \right\|_2 \sum_{\ell \geq 2} \left\| \begin{bmatrix} \mathbf{x}_{S_\ell} \\ \mathbf{e}_{T_\ell} \end{bmatrix} \right\|_2 \right) \\
&\leq \frac{\sqrt{1 + \delta_{2J}}}{1 - \delta_{2J}} \left\| \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix} \right\|_2 \|\mathbf{Ax} + \mathbf{e}\|_2 + \frac{\delta_{3J}}{1 - \delta_{2J}} \left\| \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix} \right\|_2 \sum_{\ell \geq 2} \left\| \begin{bmatrix} \mathbf{x}_{S_\ell} \\ \mathbf{e}_{T_\ell} \end{bmatrix} \right\|_2.
\end{aligned}$$

Dividing both sides by $\left\| \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix} \right\|_2$ and using the fact that $\delta_{2J} \leq \delta_{3J}$ yields

$$\left\| \begin{bmatrix} \mathbf{x}_S \\ \mathbf{e}_T \end{bmatrix} \right\|_2 \leq \left\| \begin{bmatrix} \mathbf{x}_{S \cup S_1} \\ \mathbf{e}_{T \cup T_1} \end{bmatrix} \right\|_2 \leq \frac{\sqrt{1 + \delta_{3J}}}{1 - \delta_{3J}} \|\mathbf{Ax} + \mathbf{e}\|_2 + \frac{\delta_{3J}}{1 - \delta_{3J}} \sum_{\ell \geq 2} \left\| \begin{bmatrix} \mathbf{x}_{S_\ell} \\ \mathbf{e}_{T_\ell} \end{bmatrix} \right\|_2.$$

Thanks to the definition of the S_ℓ 's, for every $j \in S_\ell$ and $k \in S_{\ell-1}$ with $\ell \geq 2$, we have $|x_j|u_j^{-1} \leq |x_k|u_k^{-1}$. Moreover, let $\alpha_k := u_k^2/|S_{\ell-1}|u$. Then, $\sum_{k \in S_{\ell-1}} \alpha_k = 1$ and, for every $j \in S_\ell$ and $\ell \geq 2$, we see that

$$\frac{|x_j|}{u_j} \leq \sum_{k \in S_{\ell-1}} \alpha_k \frac{|x_k|}{u_k} = \frac{\|\mathbf{x}_{S_{\ell-1}}\|_{1,u}}{|S_{\ell-1}|u} \leq \frac{\|\mathbf{x}_{S_{\ell-1}}\|_{1,u}}{K - \|\mathbf{u}\|_\infty^2}.$$

Therefore, for any $\ell \geq 2$, using that $K \geq \frac{4}{3}\|\mathbf{u}\|_\infty^2$, we have

$$\|\mathbf{x}_{S_\ell}\|_2 \leq \sqrt{K} \frac{\|\mathbf{x}_{S_{\ell-1}}\|_{1,u}}{K - \|\mathbf{u}\|_\infty^2} \leq \frac{4\|\mathbf{x}_{S_{\ell-1}}\|_{1,u}}{\sqrt{K}}.$$

Using an analogous argument, we obtain $\|\mathbf{e}_{T_\ell}\|_2 \leq \frac{4\|\mathbf{e}_{T_{\ell-1}}\|_{1,v}}{\sqrt{H}}$ for every $\ell \geq 2$.

Therefore, we estimate

$$\begin{aligned}
\left\| \begin{bmatrix} \mathbf{x}_S \\ \mathbf{e}_T \end{bmatrix} \right\|_2 &\leq \frac{\delta_{3J}}{1 - \delta_{3J}} \sum_{\ell \geq 2} \left\| \begin{bmatrix} \mathbf{x}_{S_\ell} \\ \mathbf{e}_{T_\ell} \end{bmatrix} \right\|_2 + \frac{\sqrt{1 + \delta_{3J}}}{1 - \delta_{3J}} \|\mathbf{Ax} + \mathbf{e}\|_2 \\
&\leq \frac{\delta_{3J}}{1 - \delta_{3J}} \sum_{\ell \geq 2} (\|\mathbf{x}_{S_\ell}\|_2 + \|\mathbf{e}_{T_\ell}\|_2) + \frac{\sqrt{1 + \delta_{3J}}}{1 - \delta_{3J}} \|\mathbf{Ax} + \mathbf{e}\|_2 \\
&\leq \frac{4\delta_{3J}}{1 - \delta_{3J}} \sum_{\ell \geq 2} \left(\frac{\|\mathbf{x}_{S_{\ell-1}}\|_{1,u}}{\sqrt{K}} + \lambda \frac{\|\mathbf{e}_{T_{\ell-1}}\|_{1,v}}{\lambda\sqrt{H}} \right) + \frac{\sqrt{1 + \delta_{3J}}}{1 - \delta_{3J}} \|\mathbf{Ax} + \mathbf{e}\|_2 \\
&= \frac{4\delta_{3J}}{1 - \delta_{3J}} \left(\frac{\|\mathbf{x}_{S^c}\|_{1,u}}{\sqrt{K}} + \frac{\lambda\|\mathbf{e}_{T^c}\|_{1,v}}{\lambda\sqrt{H}} \right) + \frac{\sqrt{1 + \delta_{3J}}}{1 - \delta_{3J}} \|\mathbf{Ax} + \mathbf{e}\|_2 \\
&\leq \frac{4\delta_{3J}}{1 - \delta_{3J}} \max \left\{ \frac{1}{\sqrt{K}}, \frac{1}{\lambda\sqrt{H}} \right\} \left\| \begin{bmatrix} \mathbf{x}_{S^c} \\ \mathbf{e}_{T^c} \end{bmatrix} \right\|_{1,w_\lambda} + \frac{\sqrt{1 + \delta_{3J}}}{1 - \delta_{3J}} \left\| M \begin{bmatrix} \mathbf{x} \\ \mathbf{e} \end{bmatrix} \right\|_2.
\end{aligned}$$

Now, observing that

$$\max \left\{ \frac{1}{\sqrt{K}}, \frac{1}{\lambda\sqrt{H}} \right\} = \frac{\Theta}{\sqrt{H + \lambda^2 K}},$$

we have shown the 2-level null-space property with constants ρ and τ defined by (68). For ρ to be strictly smaller than 1, we need condition (67). This concludes the proof. \square

Remark 5.22. The constant $4/3$ in the hypothesis $\mathbf{J} \geq \frac{4}{3}(\|\mathbf{u}\|_\infty^2, \|\mathbf{v}\|_\infty^2)$ of Theorem 5.21 can be made arbitrarily close to 1. This particular choice has been made in view of Lemma 5.7.

Finally, we establish the 2-level weighted restricted isometry property for the augmented matrix $[\mathbf{A} \quad \mathbf{I}]$ with high probability provided a suitable lower bound for the sample complexity m .

Theorem 5.23 (Sample complexity \Rightarrow 2-level weighted restricted isometry property). *Let $\mathbf{J} = (K(s), H)$, $0 < \delta < 1$, and suppose that*

$$m \gtrsim K(s) \max\{L(s, n, \delta, \varepsilon), \delta^{-2}H\}$$

Then, with probability at least $1 - \varepsilon$, the matrix $[\mathbf{A} \quad \mathbf{I}]$ has the 2-level weighted restricted isometry property of order \mathbf{J} with weights $(\mathbf{u}_\Lambda, \mathbf{v})$ and constant $\delta_{\mathbf{J}} \leq \delta$.

Proof. First, notice that [3, Lemma 3.14] can be easily adapted to the weighted case by noticing that $|\langle \mathbf{A}\mathbf{x}, \mathbf{e} \rangle| \leq \sqrt{\frac{K(s)H}{m}} \|\mathbf{x}\|_2 \|\mathbf{e}\|_2$, due to the assumption $\mathbf{v} \geq 1$. As a consequence, [3, Theorem 3.15] admits an analogous generalization to the weighted case. The thesis now follows from Theorem 5.5. \square

Finally, we introduce the concept of quasi-best weighted K -term approximation, introduced and studied in [51].

Definition 5.24 (Quasi-best weighted K -term approximation). Given $\mathbf{z} \in \mathbb{C}^n$, the corresponding *weighted quasi-best K -term approximation error* is defined as

$$\tilde{\sigma}_K(\mathbf{z})_{p, \mathbf{w}} = \|\mathbf{z} - \mathbf{z}_S\|_{p, \mathbf{w}},$$

where $S = \{\pi(1), \dots, \pi(s)\}$ where s is the maximal number such that $\sum_{j=1}^s w_{\pi(j)}^2 \leq K$, and π is some permutation such that $|z_{\pi(j)}|^p w_j^{-p} \geq |z_{\pi(j+1)}|^p w_{j+1}^{-p}$ for every $j \in [n]$.

Equipped with the technical elements discussed above, we are now ready to prove the robustness of WLAD-LASSO under unknown error.

5.2.2 Robustness of WLAD-LASSO

We prove robust recovery error estimates for the WLAD-LASSO decoder (60) under unknown error.

Theorem 5.25 (WLAD-LASSO recovery estimate). *Let $0 < \varepsilon < 1$, $2 \leq s \leq 2^{d+1}$, $H > 0$, $\Lambda = \Lambda_{d,s}^{HC}$ be the hyperbolic cross index set defined in (16), and $\{\phi_{\mathbf{i}}\}_{\mathbf{i} \in \mathbb{N}_0^d}$ be the tensor Legendre or Chebyshev polynomial basis. Let $0 < \delta < 1$ be such that*

$$\delta \leq \frac{1}{1 + 4\Theta}, \quad \text{with } \Theta = \frac{\sqrt{K(s) + \lambda^2 H}}{\min\{\sqrt{K(s)}, \lambda\sqrt{H}\}}.$$

Draw $\mathbf{t}_1, \dots, \mathbf{t}_m$ independently according to the corresponding measure ν , with

$$m \gtrsim s^{\gamma/2} \max\{L(s, n, \delta, \varepsilon), \delta^{-2}H\} \tag{69}$$

where γ is defined as in (20) and $L(s, n, \delta, \varepsilon)$ is the polylogarithmic factor defined in (39). Then, the following holds with probability at least $1 - \varepsilon$. For any $f \in L_\nu^2(D) \cap L^\infty(D)$ expanded as in (3), the approximation \tilde{f} defined as in (13) computed using the WLAD-LASSO decoder (60) satisfies

$$\|f - \tilde{f}\|_{L^\infty} + \lambda \|\mathbf{e} - (\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_\Lambda)\|_{1, \mathbf{v}} \leq C_1 (\sigma_{s,L}(\mathbf{x})_{1, \mathbf{u}} + \lambda \sigma_H(\mathbf{e})_{1, \mathbf{v}}). \tag{70}$$

and, additionally assuming $H \geq 2\|\mathbf{v}\|_\infty^2$,

$$\|f - \tilde{f}\|_{L_v^2} + \|\mathbf{e} - (\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_\Lambda)\|_2 \leq C_2(1 + \sqrt{\Theta}) \left(\frac{\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}}}{s^{\gamma/2}} + \frac{\sigma_H(\mathbf{e})_{1,\mathbf{v}}}{\sqrt{H}} \right) + \|f - f_\Lambda\|_{L_v^2}, \quad (71)$$

where f_Λ is defined as in (4) and the constants are

$$C_1 = \frac{3 + \rho}{1 - \rho}, \quad C_2 = \frac{4\sqrt{2}(1 + \rho) \max\{\rho, 2\sqrt{\rho}\}}{1 - \rho},$$

where $\rho = 4\delta/(1 - \delta)$, $\tau = \sqrt{1 - \delta}/(1 - \delta)$.

Proof. First, we observe that the sample complexity lower bound (69) is sufficient to guarantee the 2-level weighted restricted isometry property and the 2-level weighted robust null space property for the matrix $[\mathbf{A} \quad \mathbf{I}]$ with probability at least $1 - \varepsilon$, due to Theorems 5.23 & 5.21.

Now, considering the augmented formulation (63) of WLAD-LASSO, we apply Theorem 5.18 with $\mathbf{M} = [\mathbf{A} \quad \mathbf{I}]$, $\mathbf{z} = \begin{bmatrix} \mathbf{x}_\Lambda \\ \mathbf{e} \end{bmatrix}$, and $\hat{\mathbf{z}} = \begin{bmatrix} \hat{\mathbf{x}}_\Lambda \\ \hat{\mathbf{e}} \end{bmatrix}$. We obtain

$$\|\mathbf{x}_\Lambda - \hat{\mathbf{x}}_\Lambda\|_{1,\mathbf{u}} + \lambda\|\mathbf{e} - \hat{\mathbf{e}}\|_{1,\mathbf{v}} \leq \frac{2(1 + \rho)}{1 - \rho} (\sigma_{s,L}(\mathbf{x})_{1,\mathbf{u}} + \lambda\sigma_H(\mathbf{e})_{1,\mathbf{v}}), \quad (72)$$

which, for reasons analogous to those of the previous robust recovery theorems, implies (70).⁶

As a second step, we now convert the weighted ℓ^1 error estimate (72) into an L_v^2 error estimate as usual. This will require some extra work in this case.

Recalling Definition 5.24, let $\mathbf{d} := \mathbf{x}_\Lambda - \hat{\mathbf{x}}_\Lambda$ and $\mathbf{f} := \mathbf{e} - \hat{\mathbf{e}}$. Let S and T be such that $\|\mathbf{d} - \mathbf{d}_S\|_{1,\mathbf{u}} = \|\mathbf{d}_{S^c}\|_{1,\mathbf{u}} = \tilde{\sigma}_{K(s)}(\mathbf{d})_{1,\mathbf{u}}$ and $\|\mathbf{f} - \mathbf{f}_S\|_{1,\mathbf{v}} = \|\mathbf{f}_{T^c}\|_{1,\mathbf{v}} = \tilde{\sigma}_H(\mathbf{f})_{1,\mathbf{v}}$. Moreover, using the 2-level notation of Definition 5.15, define $\mathbf{g} := \begin{bmatrix} \mathbf{d} \\ \mathbf{f} \end{bmatrix}$ and $\mathbf{J} := (K(s), H)$. Recalling (61), let $\mathbf{w}_\lambda := \begin{bmatrix} \mathbf{u} \\ \lambda\mathbf{v} \end{bmatrix}$, $J_\lambda := K(s) + \lambda^2 H$. Moreover, define

$$\theta_d := \min_{i \in S} \frac{|d_i|}{u_i}, \quad \theta_f := \min_{j \in T} \frac{|f_j|}{v_j}, \quad \theta := \max \left\{ \theta_d, \frac{\theta_f}{\lambda} \right\}.$$

Using the fact that S and T realize the quasi-best approximation error, we estimate

$$\|\mathbf{g}_{(S \cup T)^c}\|_2^2 = \sum_{i \notin S} |d_i|^2 + \sum_{j \notin T} |f_j|^2 \leq \theta_d \sum_{i \notin S} u_i |d_i| + \theta_f \sum_{j \notin T} v_j |f_j| \quad (73)$$

$$\leq \theta \left(\sum_{i \notin S} u_i |d_i| + \lambda \sum_{j \notin T} v_j |f_j| \right) = \theta \|\mathbf{g}_{(S \cup T)^c}\|_{1,\mathbf{w}_\lambda} \quad (74)$$

Similarly to the proof of Theorem 5.21, let $\alpha_k := u_k^2/|S|_{\mathbf{u}}$. Notice that $\alpha_k \leq (K(s) - \|\mathbf{u}_\Lambda\|_\infty^2)^{-1} u_k^2$ (due to the maximality of S , realizing the quasi-best term approximation) and that $\sum_{k \in S} \alpha_k = 1$. Now we observe that

$$\theta_d^2 = \sum_{i \in S} \alpha_i \min_{k \in S} \left(\frac{|d_k|}{u_k} \right)^2 \leq \sum_{i \in S} \alpha_i \left(\frac{|d_i|}{u_i} \right)^2 \leq \frac{1}{K(s) - \|\mathbf{u}_\Lambda\|_\infty^2} \sum_{i \in S} u_i^2 \left(\frac{|d_i|}{u_i} \right)^2 = \frac{1}{K(s) - \|\mathbf{u}_\Lambda\|_\infty^2} \|\mathbf{d}_S\|_2^2.$$

Recalling that $\frac{3}{4}K(s) \geq \|\mathbf{u}_\Lambda\|_\infty^2$ (due to Lemma 5.7), we get $\theta_d \leq \frac{2}{\sqrt{K(s)}} \|\mathbf{d}_S\|_2$. Similarly, we can deduce that $\theta_f \leq \frac{\sqrt{2}}{\sqrt{H}} \|\mathbf{f}_T\|_2$ thanks to the hypothesis $H \geq 2\|\mathbf{v}\|_\infty^2$. Therefore, using the 2-level weighted robust null space property, we obtain

$$\theta = \max \left\{ \theta_d, \frac{\theta_f}{\lambda} \right\} \leq \frac{2(\|\mathbf{d}_S\|_2 + \|\mathbf{f}_T\|_2)}{\min \{\sqrt{K(s)}, \lambda\sqrt{H}\}} \leq \frac{2\sqrt{2}\|\mathbf{g}_{S \cup T}\|_2}{\min \{\sqrt{K(s)}, \lambda\sqrt{H}\}} \leq \frac{2\sqrt{2}\rho\|\mathbf{g}_{(S \cup T)^c}\|_{1,\mathbf{w}_\lambda}}{\min \{\sqrt{K(s)}, \lambda\sqrt{H}\}\sqrt{J_\lambda}}.$$

⁶Notice that since $\mathbf{M}(\mathbf{z} - \hat{\mathbf{z}}) = \mathbf{0}$ the constant τ does not appear in (72). Indeed, we are just using a 2-level weighted version of the so-called *stable* null space property (see [31, §4.2]).

Plugging the above inequality into (74) and defining $\beta := \frac{\|\mathbf{g}_{(S \cup T)^c}\|_{1, \mathbf{w}_\lambda}}{\sqrt{J_\lambda}}$, we deduce that

$$\|\mathbf{g}_{(S \cup T)^c}\|_2^2 \leq \theta \|\mathbf{g}_{(S \cup T)^c}\|_{1, \mathbf{w}_\lambda} \leq 2\sqrt{2}\rho\Theta\beta^2 \leq 4\rho\Theta\beta^2.$$

Using the above estimate and the weighted 2-level robust null space property again, we see that

$$\begin{aligned} \frac{1}{\sqrt{2}} (\|\mathbf{d}\|_2 + \|\mathbf{f}\|_2) &\leq \|\mathbf{g}\|_2 \leq \|\mathbf{g}_{S \cup T}\|_2 + \|\mathbf{g}_{(S \cup T)^c}\|_2 \\ &\leq \frac{\rho}{\sqrt{J_\lambda}} \|\mathbf{g}_{(S \cup T)^c}\|_{1, \mathbf{w}_\lambda} + \|\mathbf{g}_{(S \cup T)^c}\|_2 \\ &\leq \rho\beta + 2\sqrt{\rho\Theta}\beta \leq \max\{\rho, 2\sqrt{\rho}\} (1 + \sqrt{\Theta}) \beta. \end{aligned}$$

Now, taking advantage of (72), we have

$$\beta \leq \frac{\|\mathbf{g}\|_{1, \mathbf{w}}}{\sqrt{J_\lambda}} \leq \frac{2(1+\rho)}{1-\rho} \frac{\sigma_{s,L}(\mathbf{x})_{1, \mathbf{u}} + \lambda \sigma_H(\mathbf{e})_{1, \mathbf{v}}}{\sqrt{K(s) + \lambda^2 H}} \leq \frac{2(1+\rho)}{1-\rho} \left(\frac{\sigma_{s,L}(\mathbf{x})_{1, \mathbf{u}}}{\sqrt{K(s)}} + \lambda \frac{\sigma_H(\mathbf{e})_{1, \mathbf{v}}}{\sqrt{H}} \right),$$

The above inequalities, combined with Lemma 2.2 and with the inequality $\|f - \tilde{f}\|_{L_v^2} \leq \|f - f_\Lambda\|_{L_v^2} + \|\mathbf{d}\|_2$ yield (71). This concludes the proof. \square

6 Conclusions

In the context of sparse high-dimensional approximation from pointwise samples, we have considered four decoders for weighted ℓ^1 minimization: weighted quadratically-constrained basis pursuit, weighted LASSO, weighted square-root LASSO, and weighted LAD-LASSO. We have compared these decoders from the theoretical and the numerical perspectives, focusing on the case where the samples are corrupted by unknown error (such as truncation, discretization, and numerical error).

On the theoretical side, we have proved uniform robust recovery guarantees for all the decoders considered, showing that they achieve the best s -term approximation error in lower sets, provided a suitable lower bound on the sample complexity and up to the error level (see §2.4 and Theorems 5.10, 5.13, 5.14, and 5.25). Our analysis suggests optimal strategies for the choice of the respective tuning parameters (see equations (24), (25), (28), and (31)).

From the numerical viewpoint, we have compared the decoders' performance on a synthetic example, where f is explicitly available (§4.2 & 4.3), and on a more applicable case, where the function is defined as the quantity of interest of a parametric ODE or PDE with random inputs (§4.4). The first set of experiments corroborates the robustness and the reliability of the decoders under examination and confirms the optimality of the choices of the tuning parameters suggested by the theory. Conversely, for the latter examples, our numerical results suggest that careful parameter tuning may not always be needed, possibly due to the structured nature of the error arising in these types of problems (see §4.4.3). A proper theoretical understanding of this phenomenon is still an open issue. Moreover, devising suitable strategies to incorporate some *a priori* information on f to the proposed approximation scheme is another open problem. For example, when f is known to have more variability along certain axial directions than others, one might replace the hyperbolic cross with some anisotropic set Λ or modify the weights \mathbf{u} used in the decoders. Analyzing these issues is beyond the scope of this paper and it is left to future work.

Finally, we remark the following. In light of the comparison made in this paper, we suggest that the square-root LASSO should be the decoder of choice out of those studied (apart from the case of sparse corruptions, where the best choice is of course the weighted LAD-LASSO). Indeed, this decoder has the same theoretical guarantees as the other ones and exhibits comparable (and sometimes better) numerical performance. Yet, the optimal choice of its tuning parameter does not depend on the unknown error. Although this fact is well-known in the statistics community, we think that it has not been fully exploited in high-dimensional function approximation and in compressed sensing. We hope that this paper could convince the reader to consider the square-root LASSO as a valid alternative to quadratically-constrained basis pursuit and LASSO in their future investigations.

Acknowledgements

BA, AB and SB acknowledge the Natural Sciences and Engineering Research Council of Canada through grant 611675 and the Alfred P. Sloan Foundation and the Pacific Institute for the Mathematical Sciences (PIMS) Collaborative Research Group “High-Dimensional Data Analysis”. SB acknowledges the support of the PIMS Post-doctoral Training Center in Stochastics. The authors are grateful to Claire Boyer, John Jakeman, Richard Lockhart, Akil Narayan, and Clayton G. Webster for interesting and fruitful discussions.

References

- [1] B. Adcock. Infinite-dimensional compressed sensing and function interpolation. *Found. Comput. Math.*, to appear, 2017.
- [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, A. Narayan, and U. Author. Compressed sensing with sparse corruptions: Fault-tolerant sparse collocation approximations. *arXiv:1703.00135*, 2017.
- [4] B. Adcock, S. Brugiapaglia, and C. G. Webster. Compressed sensing approaches for polynomial approximation of high-dimensional functions. *arXiv:1703.06987*, 2017.
- [5] B. Adcock, A. C. Hansen, C. Poon, and B. Roman. Breaking the coherence barrier: A new theory for compressed sensing. In *Forum of Mathematics, Sigma*, volume 5. Cambridge University Press, 2017.
- [6] S. Arlot and A. Celisse. A survey of cross-validation procedures for model selection. *Statistics surveys*, 4:40–79, 2010.
- [7] O. Arslan. Weighted LAD-LASSO method for robust parameter estimation and variable selection in regression. *Comput. Statist. Data Anal.*, 56(6):1952–1965, 2012.
- [8] P. Babu and P. Stoica. Connection between spice and square-root lasso for sparse parameter estimation. *Signal Process.*, 95:10–14, 2014.
- [9] J. Bäck, F. Nobile, L. Tamellini, and R. Tempone. Stochastic spectral galerkin and collocation methods for pdes with random coefficients: A numerical comparison. In J. S. Hesthaven and E. M. Rønquist, editors, *Spectral and High Order Methods for Partial Differential Equations: Selected papers from the ICOSAHOM ’09 conference, June 22–26, Trondheim, Norway*, pages 43–62. Springer Berlin Heidelberg, 2011.
- [10] J. Ballani and L. Grasedyck. Hierarchical tensor approximation of output quantities of parameter-dependent pdes. *SIAM/ASA J. Uncertainty Quantification*, 3(1):852–872, 2015.
- [11] A. Bastounis and A. C. Hansen. On the absence of the RIP in real-world applications of compressed sensing and the RIP in levels. *arXiv:1411.4449*, 2014.
- [12] A. Belloni, V. Chernozhukov, and L. Wang. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011.
- [13] A. Belloni, V. Chernozhukov, and L. Wang. Pivotal estimation via square-root lasso in nonparametric regression. *Ann. Statist.*, 42(2):757–788, 2014.
- [14] P. G. Bridges, K. B. Ferreira, M. A. Heroux, and M. Hoemmen. Fault-tolerant linear solvers via selective reliability. *arXiv:1206.1390*, 2012.
- [15] S. Brugiapaglia and B. Adcock. Robustness to unknown error in sparse regularization. *arXiv:1705.10299*, 2017.
- [16] S. Brugiapaglia, B. Adcock, and R. K. Archibald. Recovery guarantees for compressed sensing with unknown errors. In *Sampling Theory and Applications (SampTA), 2017 International Conference on*. IEEE, 2017.
- [17] F. Bunea, J. Lederer, and Y. She. The group square-root lasso: Theoretical properties and fast algorithms. *IEEE Trans. Inform. Theory*, 60(2):1313–1325, 2014.
- [18] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52(2):489–509, 2006.

- [19] E. J. Candès, M. B. Wakin, and S. P. Boyd. Enhancing sparsity by reweighted ℓ^1 minimization. *J. Fourier Anal. Appl.*, 14(5):877–905, 2008.
- [20] A. Chkifa, A. Cohen, G. Migliorati, F. Nobile, and 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.
- [21] A. Chkifa, A. Cohen, and C. Schwab. High-Dimensional Adaptive Sparse Polynomial Interpolation and Applications to Parametric PDEs. *Found. Comput. Math.*, 14(4):601–633, Aug 2014.
- [22] A. Chkifa, A. Cohen, and C. Schwab. Breaking the curse of dimensionality in sparse polynomial approximation of parametric PDEs. *J. Math. Pures Appl.*, 103(2):400–428, 2015.
- [23] A. Chkifa, N. Dexter, H. Tran, and C. G. Webster. Polynomial approximation via compressed sensing of high-dimensional functions on lower sets. *Math. Comp.*, to appear, 2017.
- [24] A. Cohen, R. DeVore, and C. Schwab. Convergence Rates of Best N-term Galerkin Approximations for a Class of Elliptic sPDEs. *Found. Comput. Math.*, 10(6):615–646, Dec 2010.
- [25] A. Cohen, R. DeVore, and C. Schwab. Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE’s. *Anal. Appl.*, 9(01):11–47, 2011.
- [26] C. de Boor and A. Ron. On multivariate polynomial interpolation. *Constr. Approx.*, 6(3):287–302, Sep 1990.
- [27] D. L. Donoho. Compressed Sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.
- [28] D. L. Donoho and B. F. Logan. Signal Recovery and the Large Sieve. *SIAM J. Appl. Math.*, 52(2):577–591, 1992.
- [29] A. Doostan and H. Owhadi. A non-adapted sparse approximation of PDEs with stochastic inputs. *J. Comput. Phys.*, 230(8):3015–3034, 2011.
- [30] N. Dyn and M. S. Floater. Multivariate polynomial interpolation on lower sets. *J. Approx. Theory*, 177(Supplement C):34 – 42, 2014.
- [31] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Appl. Numer. Harmon. Anal. Springer Science+Business Media, New York, 2013.
- [32] M. P. Friedlander, H. Mansour, R. Saab, and O. Yilmaz. Recovering compressively sampled signals using partial support information. *IEEE Trans. Inform. Theory*, 58(2):1122–1134, 2012.
- [33] X. Gao. *Penalized methods for high-dimensional least absolute deviations regression*. PhD thesis, The University of Iowa, 2008.
- [34] X. Gao and J. Huang. Asymptotic analysis of high-dimensional lad regression with lasso. *Statist. Sinica*, pages 1485–1506, 2010.
- [35] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008.
- [36] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- [37] T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical Learning with Sparsity: the Lasso and Generalizations*. CRC press, 2015.
- [38] F. Hecht. New development in FreeFem++. *J. Numer. Math.*, 20(3-4):251–265, 2012.
- [39] J. Huang, S. Ma, and C.-H. Zhang. Adaptive lasso for sparse high-dimensional regression models. *Statist. Sinica*, pages 1603–1618, 2008.
- [40] J.D. Jakeman, M.S. Eldred, and K. Sargsyan. Enhancing ℓ^1 -minimization estimates of polynomial chaos expansions using basis selection. *J. Comput. Phys.*, 289:18 – 34, 2015.
- [41] J. N. Laska, M. A. Davenport, and R. G. Baraniuk. Exact signal recovery from sparsely corrupted measurements through the pursuit of justice. In *Signals, Systems and Computers, 2009 Conference Record of the Forty-Third Asilomar Conference on*, pages 1556–1560. IEEE, 2009.
- [42] Q. Li and L. Wang. Robust change point detection method via adaptive lad-lasso. *Statist. Papers*, pages 1–13, 2017.

- [43] X. Li. Compressed sensing and matrix completion with constant proportion of corruptions. *Constr. Approx.*, 37(1):73–99, 2013.
- [44] B. F. Logan. *Properties of high-pass signals*. PhD thesis, Columbia University, 1965.
- [45] G. G. Lorentz and R. A. Lorentz. Solvability problems of bivariate interpolation i. *Constr. Approx.*, 2(1):153–169, Dec 1986.
- [46] G. Migliorati, F. Nobile, E. von Schwerin, and R. Tempone. Analysis of Discrete L^2 Projection on Polynomial Spaces with Random Evaluations. *Found. Comput. Math.*, 14(3):419–456, Jun 2014.
- [47] N. H. Nguyen and T. D. Tran. Exact recoverability from dense corrupted observations via ℓ_1 -minimization. *IEEE Trans. Inform. Theory*, 59(4):2017–2035, 2013.
- [48] J. Peng, J. Hampton, and A. Doostan. A weighted ℓ_1 minimization approach for sparse polynomial chaos expansions. *J. Comput. Phys.*, 267:92–111, 2014.
- [49] V. Pham and L. El Ghaoui. Robust sketching for multiple square-root lasso problems. In *Artificial Intelligence Stat.*, pages 753–761, 2015.
- [50] H. Rauhut and C. Schwab. Compressive sensing Petrov-Galerkin approximation of high-dimensional parametric operator equations. *Math. Comp.*, 86(304):661–700, 2017.
- [51] H. Rauhut and R. Ward. Interpolation via weighted ℓ_1 minimization. *Appl. Comput. Harmon. Anal.*, 40(2):321–351, 2016.
- [52] Y. Shin and D. Xiu. Correcting data corruption errors for multivariate function approximation. *SIAM J. Sci. Comput.*, 38(4):A2492–A2511, 2016.
- [53] L. Stankovic, S. Stankovic, and M. Amin. Missing samples analysis in signals for applications to l-estimation and compressive sensing. *Signal Process.*, 94:401–408, 2014.
- [54] B. Stucky and S. A. van de Geer. Sharp oracle inequalities for square root regularization. *arXiv:1509.04093*, 2015.
- [55] C. Studer, P. Kuppinger, G. Pope, and H. Bolcskei. Recovery of sparsely corrupted signals. *IEEE Trans. Inform. Theory*, 58(5):3115–3130, 2012.
- [56] D. Su. Compressed sensing with corrupted fourier measurements. *arXiv:1607.04926*, 2016.
- [57] D. Su. Data recovery from corrupted observations via l1 minimization. *arXiv:1601.06011*, 2016.
- [58] T. Sun and C.-H. Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012.
- [59] X. Tian, J. R. Loftus, and J. E. Taylor. Selective inference with unknown variance via the square-root lasso. *arXiv:1504.08031*, 2015.
- [60] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, pages 267–288, 1996.
- [61] S. A. van de Geer. *Estimation and testing under sparsity*. Springer, 2016.
- [62] H. Wang, G. Li, and G. Jiang. Robust regression shrinkage and consistent variable selection through the LAD-Lasso. *J. Bus. Econom. Statist.*, 25(3):347–355, 2007.
- [63] J. Wright and Y. Ma. Dense error correction via ℓ^1 -minimization. *IEEE Trans. Inform. Theory*, 56(7):3540–3560, 2010.
- [64] J. Xu. *Parameter estimation, model selection and inferences in L1-based linear regression*. PhD thesis, Columbia University, 2005.
- [65] J. Xu and Z. Ying. Simultaneous estimation and variable selection in median regression using lasso-type penalty. *Ann. Inst. Statist. Math.*, 62(3):487–514, 2010.
- [66] L. Yan, L. Guo, and D. Xiu. Stochastic collocation algorithms using ℓ_1 -minimization. *Int. J. Uncertain. Quantif.*, 2(3):279–293, 2012.
- [67] X. Yang and G. E. Karniadakis. Reweighted ℓ^1 minimization method for stochastic elliptic differential equations. *J. Comput. Phys.*, 248:87–108, 2013.
- [68] X. Yu and S. J. Baek. Sufficient conditions on stable recovery of sparse signals with partial support information. *IEEE Signal Process. Lett.*, 20(5):539–542, 2013.
- [69] H. Zou. The Adaptive Lasso and its Oracle Properties. *J. Amer. Statist. Assoc.*, 101(476):1418–1429, 2006.