



A Unified Approach to Uniform Signal Recovery From Nonlinear Observations

Martin Genzel¹ · Alexander Stollenwerk²

Received: 26 November 2020 / Revised: 16 August 2021 / Accepted: 5 January 2022 /
Published online: 30 March 2022

© SFoCM 2022

Abstract

Recent advances in quantized compressed sensing and high-dimensional estimation have shown that signal recovery is even feasible under strong nonlinear distortions in the observation process. An important characteristic of associated guarantees is uniformity, i.e., recovery succeeds for an entire class of structured signals with a fixed measurement ensemble. However, despite significant results in various special cases, a general understanding of uniform recovery from nonlinear observations is still missing. This paper develops a unified approach to this problem under the assumption of i.i.d. sub-Gaussian measurement vectors. Our main result shows that a simple least-squares estimator with any convex constraint can serve as a universal recovery strategy, which is outlier robust and does not require explicit knowledge of the underlying nonlinearity. Based on empirical process theory, a key technical novelty is an approximative increment condition that can be implemented for all common types of nonlinear models. This flexibility allows us to apply our approach to a variety of problems in nonlinear compressed sensing and high-dimensional statistics, leading to several new and improved guarantees. Each of these applications is accompanied by a conceptually simple and systematic proof, which does not rely on any deeper properties of the observation model. On the other hand, known local stability properties

Communicated by Thomas Strohmer.

MG acknowledges support by the DFG Priority Programme DFG-SPP 1798 Grant DI 2120/1-1. AS acknowledges support by the Fonds de la Recherche Scientifique – FNRS under Grant n° T.0136.20 (Learn2Sense).

✉ Martin Genzel
m.genzel@uu.nl

Alexander Stollenwerk
alexander.stollenwerk@uclouvain.be

¹ Utrecht University, Mathematical Institute, Budapestlaan 6, 3584, CD Utrecht, Netherlands

² UCLouvain, ICTEAM Institute, Avenue Georges Lemaître 4-6/L4, B1348 Louvain-la-Neuve, Belgium

can be incorporated into our framework in a plug-and-play manner, thereby implying near-optimal error bounds.

Keywords Uniform recovery · High-dimensional estimation · Nonlinear observations · Quantized compressed sensing · Empirical processes

Mathematics Subject Classification 62F30 · 62F35 · 62J07 · 94A12

1 Introduction

This paper is concerned with the following fundamental reconstruction task:

Problem 1 Consider a set of signals $\mathcal{X} \subset \mathbb{R}^p$ and let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be a collection of measurement vectors. Moreover, let $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ be a scalar output function. Under what conditions can the following recovery problem be solved *uniformly* for all $\mathbf{x} \in \mathcal{X}$: Assume that \mathbf{x} is observed in the form

$$y_i := F(\mathbf{a}_i, \mathbf{x}) + v_i, \quad i = 1, \dots, m, \quad (1.1)$$

where $v_1, \dots, v_m \in \mathbb{R}$ is scalar noise. Given $\{(\mathbf{a}_i, y_i)\}_{i=1}^m$, is it possible to recover the underlying signal \mathbf{x} efficiently?

The prototypical instance of Problem 1 is the approach of *compressed sensing*. Dating back to the seminal works of Candès, Romberg, Tao, and Donoho [10, 11, 19], the traditional setup of compressed sensing focuses on the case of *noisy linear observations*, i.e., we have $y_i = \langle \mathbf{a}_i, \mathbf{x} \rangle + v_i$ for $i = 1, \dots, m$. In this regime, Problem 1 is nowadays fairly well understood, underpinned by various real-world applications, efficient algorithmic methods, and a rich theoretical foundation; see [21] for a comprehensive overview. In a nutshell, compressed sensing has proven that signal recovery is still feasible when $m \ll p$, supposed that the signal set \mathcal{X} carries some low-dimensional structure, e.g., a variant of *sparsity*, while the measurement ensemble $\{\mathbf{a}_i\}_{i=1}^m$ follows an appropriate random design. The *uniformity* over \mathcal{X} plays an important role in this context, since the measurement device—determined by $\{\mathbf{a}_i\}_{i=1}^m$ in our case—is typically fixed in applications and should allow for the reconstruction of *all* (or most) signals in \mathcal{X} . But also apart from this practical relevance, the above quest for uniform recovery is an interesting mathematical problem in its own right. It is significantly more involved than its non-uniform counterpart, where the $\{\mathbf{a}_i\}_{i=1}^m$ may vary for each $\mathbf{x} \in \mathcal{X}$.

The present work is devoted to Problem 1 in the much less explored situation of *nonlinear* output functions. In fact, many conceptions from compressed sensing theory, such as the restricted isometry or nullspace property, are tailored to linear models and do not carry over directly to the nonlinear case. As we will see in the next subsection, the presence of nonlinear measuring components is not merely an academic concern but affects many problems in signal processing and high-dimensional statistics.

1.1 Prior Art

There exist two branches of research that are particularly relevant to this work. The first one is based in the field of (memoryless) *quantized compressed sensing*, which deals with the fact that analog (linear) measurements often need to be quantized before further processing in practice. A common scenario in this respect is *1-bit compressed sensing* where only a single bit of information is retained, e.g., if (1.1) renders observations of the form $y_i = \text{sign}(\langle \mathbf{a}_i, \mathbf{x} \rangle)$. Due to this considerable loss of information, it may come as a surprise that tractable recovery methods are still available and Problem 1 is relatively well understood in this situation. A solid theoretical basis as well as efficient algorithms have been developed over the last few years, including significant progress on lower bounds [18, 34], robustness [17, 49], advanced quantization schemes [3, 36, 37, 44, 65], and non-Gaussian measurements [1, 16–18]. While this list of references is certainly incomplete, the surveys of Dirksen [14] and Boufounos et al. [6] provide nice overviews of the field of quantized compressed sensing.

The above achievements have also given rise to several important mathematical tools. One of the most notable breakthroughs are quantized embedding results for signal recovery—a highly geometric argument based on uniform, random hyperplane tessellations, e.g., see [17, 34, 45, 49, 50]. Closer to the original ideas of compressed sensing are reconstruction guarantees relying on variants of the restricted isometry property, e.g., see [16, 20, 33, 65]. However, these techniques are strongly tailored to quantized measurements and it remains unclear how they could be extended to other instances of Problem 1.

The second branch of related literature is much less restrictive with respect to the underlying observation model. It allows (1.1) to take the form $y_i = f(\langle \mathbf{a}_i, \mathbf{x} \rangle) + v_i$, where $f: \mathbb{R} \rightarrow \mathbb{R}$ can be nonlinear and random. A pioneering work on these so-called *single-index models* is the one of Plan and Vershynin [51] (inspired by ideas of Brillinger [7]), who study the *generalized Lasso* as reconstruction method:

$$\min_{\mathbf{z} \in K} \frac{1}{m} \sum_{i=1}^m (y_i - \langle \mathbf{a}_i, \mathbf{z} \rangle)^2. \tag{P_{K,y}}$$

Here, $K \subset \mathbb{R}^p$ is a *convex* constraint set, serving as an appropriate relaxation of the actual signal set \mathcal{X} and making (P_{K,y}) tractable in many situations of interest. Although this “linearization” strategy might appear very coarse, it was shown to produce satisfactory outcomes for Gaussian measurements, even when f is highly nonlinear. A key benefit of (P_{K,y}) is that it does not require any (explicit) knowledge of the observation model, which enables various applications to signal processing and statistical estimation, e.g., see [23, Chap. 3 & 4]. The article of Plan and Vershynin [51] is just an example of a whole line of research with many related and follow-up works, e.g., see [22, 28, 29, 47, 52, 56, 60, 62]; remarkably, this approach also extends to phase-retrieval-like problems where f is an even function [25, 61, 66]. For a more detailed discussion of the literature, we refer to [52, Sec. 6] and [23, Sec. 4.2].

At first sight, the aforementioned works seem to provide a general solution to Problem 1, but they lack a crucial feature, namely uniformity. Indeed, most of these results

are based on concentration inequalities over high-dimensional signal sets, exploiting their “local geometry” around a fixed $\mathbf{x} \in \mathcal{X}$. This strategy naturally leads to *non-uniform* recovery guarantees, and we are not aware of a uniform extension in that regard. Finally, we point out that the two research areas discussed above are not independent but exhibit certain overlaps in terms of their main achievements and proof techniques, for instance, see [28, 62].

1.2 Contributions and Overview

The primary objective of this work is to fill the striking gap between the lines of research discussed in the previous subsection. On the one hand, we show that uniform recovery is not only limited to linear and quantized measurement schemes but applies to a much larger family of observation models. On the other hand, it turns out that under mild assumptions, known non-uniform guarantees for single-index and related models naturally extend to the uniform case. Apart from a unification, our theoretical approach will contribute to each of these research directions through new and improved results.

On the conceptual side, we stick to the methodology suggested by the second branch of the literature from Sect. 1.1, analyzing the performance of the generalized Lasso $(P_{K,y})$. Our main result, Theorem 2 in Sect. 2.3, establishes a uniform error bound for $(P_{K,y})$ under very general conditions, including all observation models mentioned above. This finding implies that there exists a “universal” recovery strategy, which often yields satisfactory outcomes. In this context, it is worth pointing out that we have focused on $(P_{K,y})$ mainly because of its simplicity and popularity, but similar results could be shown for other reconstruction methods as well (cf. Remark 3(3)). Therefore, the present paper is a contribution to Problem 1 in the first place, whereas the estimator of choice plays a subordinate role.¹

The recovery guarantee of Theorem 2 is formulated in a quite abstract setting, which requires some technical preparation. In order to familiarize the reader with the results of this paper, we now state a special case for single-index models with a Lipschitz continuous output function and Gaussian measurements (see Sect. 8.4 for a proof). Note that parts of the notation from Sect. 1.4 is already used here; the complexity parameters $w_t(\cdot)$ and $w(\cdot)$ correspond to the common notion of (local) mean width, which is formally introduced in Definition 2.

Theorem 1 *There exist universal constants $c, C > 0$ for which the following holds.*

Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent copies of a standard Gaussian random vector $\mathbf{a} \sim N(\mathbf{0}, \mathbf{I}_p)$. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be γ -Lipschitz, and for $g \sim N(0, 1)$, we set

$$\mu := \mathbb{E}[f(g)g] \quad \text{and} \quad r := \|f(g) - \mu g\|_{\psi_2}.$$

¹ In particular, we will not discuss any details about possible algorithmic implementations of $(P_{K,y})$, which is an important subject in its own right.

Moreover, let $\mathcal{X} \subset \mathbb{S}^{p-1}$ and let $K \subset \mathbb{R}^p$ be a convex set such that $\mu\mathcal{X} \subset K$. For $u \geq 1$ and a desired reconstruction accuracy $t \geq 0$, we assume that

$$m \geq C \cdot \left((1 + t^{-2}r^2) \cdot (w_t^2(K - \mu\mathcal{X}) + u^2) + t^{-2}(\mu + \gamma)^2 \cdot w^2(\mathcal{X}) \right), \tag{1.2}$$

where $K - \mu\mathcal{X} := \{z - \mu x : z \in K, x \in \mathcal{X}\}$. Then with probability at least $1 - \exp(-cu^2)$ on the random draw of $\{\mathbf{a}_i\}_{i=1}^m$, the following holds uniformly for all $\mathbf{x} \in \mathcal{X}$: Let $\mathbf{y} = (y_1, \dots, y_m) \in \mathbb{R}^m$ be given by

$$y_i = f(\langle \mathbf{a}_i, \mathbf{x} \rangle) + v_i, \quad i = 1, \dots, m,$$

such that $(\frac{1}{m} \sum_{i=1}^m v_i^2)^{1/2} \leq \frac{t}{20}$.² Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \mu\mathbf{x}\|_2 \leq t$.

Theorem 1 can be seen as a uniform version of Plan’s and Vershynin’s main result on single-index models [51, Thm. 1.9]. Most notably, the sampling-rate condition (1.2) implies that the dependence on the desired accuracy t is the same as in the non-uniform case. For more details on the interplay between m , t , and the complexity parameters, we refer to the discussion of Theorem 2 in Sect. 2.3. To the best of our knowledge, Theorem 1 is a new result in its own right, and in particular, it is not a consequence of the approaches outlined in the previous subsection. A novel argument is required to prove such a type of recovery guarantee; see Sect. 1.3 for an overview of the major technical challenge of our work.

The main result of Theorem 2 extends Theorem 1 by several important aspects, being in line with our key achievements:

1. *Discontinuous output functions.* The Lipschitz assumption in Theorem 1 prohibits popular models like quantized observations. It turns out that an extension in that regard is subtle. To this end, we introduce an *approximative increment condition* on the observation variable, which is satisfied for many discontinuous nonlinearities; see Sect. 2.2 for more details and Sects. 4.1–4.4 for applications to nonlinear compressed sensing.
2. *Beyond Gaussian measurements.* Compared to the previous point, it is relatively straightforward to allow for i.i.d. sub-Gaussian measurement vectors in Theorem 1. However, this relaxation leads to an additional constraint term, called the *target mismatch*, whose size depends on the distribution of the measurements and the output function; see Definition 1 in Sect. 2.1.
3. *Beyond single-index models.* In Sect. 2, we study the more general situation of Problem 1, where y_i is not necessarily representable in terms of the dot product $\langle \mathbf{a}_i, \mathbf{x} \rangle$. While this requires an additional step of abstraction, it enables for more complicated nonlinear models as well as refined recovery tasks, e.g., support estimation (see Sect. 4.5). For this purpose, we introduce so-called *target functions* in

² Here, the v_i correspond to *adversarial noise*, i.e., as long as the ℓ^2 -constraint is satisfied, they could be arbitrary (even worst-case) perturbations. This model is very common in uniform recovery and compressed sensing, cf. [21]. But note that compared to statistical noise, this comes at the price of consistent estimation, i.e., the reconstruction error may not become arbitrarily small if $m \rightarrow \infty$; see Remark 1 and 4(2) for further discussion.

Sect. 2.1, transforming a signal $x \in \mathcal{X}$ in such a way that it becomes compatible with the solution vector of $(P_{K,y})$ under observations of the form (1.1). Note that in the setup of Theorem 1, this transformation simply corresponds to rescaling x by the scalar factor μ .

4. *Outlier robustness.* The ℓ^2 -noise constraint $(\frac{1}{m} \sum_{i=1}^m v_i^2)^{1/2} \leq \frac{t}{20}$ in Theorem 1 is standard for linear models and also remains useful in the situation of Lipschitz nonlinearities. However, such a condition is overly restrictive for quantized measurement schemes, where the noise is due to (few) wrong bits instead of (small) real-valued perturbations. Theorem 2 therefore involves an important relaxation in this respect, permitting a portion of “gross” outliers in the observation vector.

Based on our findings in Theorem 2, we will also address two points of independent interest:

5. *Uniform vs. non-uniform recovery.* Section 3 isolates characteristics that are specifically attributable to uniformity. While the presence of nonlinear output functions may only affect the oversampling rate (see also Sect. 4), our focus here is on the quite implicit complexity parameter $w_t(K - \mathcal{X})$, cf. (1.2). We show that it can be always bounded by two much more informative expressions, namely a “fully localized” mean width and its global counterpart (see Proposition 1). This important simplification enables us to reuse known (non-uniform) results from the literature. We illustrate for two examples (ℓ^1 -constraints and total variation) that there is no considerable gap between the uniform and non-uniform regime.
6. *Plug-and-play via embeddings.* Section 5 demonstrates how to incorporate known random embedding results into our theoretical framework. In this way, for instance, it is possible to obtain near-optimal error decay rates for $(P_{K,y})$ with quantized measurements. This achievement is due to a more general guarantee stated in Theorem 3.

The general purpose of our methodology is to enable a more systematic study of Problem 1 than before, especially for nonlinear observation models. To a certain degree, this paper promotes an alternative path to compressed sensing theory that can do without common tools like the restricted isometry property or quantized embedding results. Despite clear overlaps, each of these approaches comes along with its own strengths and (in-)accessible regimes; see Sect. 6 for a more detailed discussion.

1.3 Technical Challenges

This subsection highlights the main technical achievements of this work. Our principal procedure is in fact relatively standard: We decompose the excess loss associated with the estimator $(P_{K,y})$ into separate empirical processes and aim at suitable lower/upper bounds for each (cf. [17, 23, 25, 41, 51]); see Sect. 2.2 for an overview of this argument.

However, the quest for *uniform* recovery as prescribed by Problem 1 makes this approach a challenging endeavor. In a nutshell, the key difficulty is that the common multiplier term turns into an *empirical product process* of the following type:

$$\frac{1}{m} \sum_{i=1}^m \xi_i(\mathbf{x}) \cdot \langle \mathbf{a}_i, \mathbf{v} \rangle, \quad \mathbf{x} \in \mathcal{X}, \quad \mathbf{v} \in (K - T\mathbf{x}) \cap t\mathbb{S}^{p-1}, \quad (1.3)$$

where the first factor takes the form $\xi_i(\mathbf{x}) = \langle \mathbf{a}_i, T\mathbf{x} \rangle - F(\mathbf{a}_i, \mathbf{x})$ and $T: \mathcal{X} \rightarrow K$ is a fixed function (think of a scalar multiplication for now). Since we require a uniform upper bound for (1.3) over both \mathbf{x} and \mathbf{v} , the regularity of the stochastic process $\{\xi_i(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ plays an important role, in the sense that it should fulfill a sub-Gaussian increment condition (see (1.4)).³ Such a regularity assumption unfortunately does not even hold for standard models like 1-bit quantization $F(\mathbf{a}_i, \mathbf{x}) = \text{sign}(\langle \mathbf{a}_i, \mathbf{x} \rangle)$. Therefore, existing chaining-based results for product processes are not directly applicable to our situation. This may also explain why there is no unified theory for nonlinear observations available so far and known uniform recovery guarantees are restricted to special cases (see Sect. 1.1).

The present article provides a general remedy for this foundational problem. Our starting point is that a given output function F can be often approximated by a more regular one, say F_t , where $t > 0$ controls the approximation accuracy. As such, this strategy can yield a more benign multiplier $\xi_{t,i}(\mathbf{x}) = \langle \mathbf{a}_i, T\mathbf{x} \rangle - F_t(\mathbf{a}_i, \mathbf{x})$, but it simply shifts the original problem to the resulting approximation error $F(\mathbf{a}_i, \mathbf{x}) - F_t(\mathbf{a}_i, \mathbf{x})$. Hence, a fundamental insight in our proof (see Step 3 in Sect. 7) is that it suffices to consider the *absolute* error $\varepsilon_{t,i} = |F(\mathbf{a}_i, \mathbf{x}) - F_t(\mathbf{a}_i, \mathbf{x})|$ instead. Indeed, the process $\{\varepsilon_{t,i}(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ fulfills an increment condition in all relevant model situations that we are aware of; see Fig. 2 in Sect. 4 for an illustration of the above argument based on the sign function. With this at hand, and apart from other technicalities, we can take advantage of a recent concentration inequality due to Mendelson [42, Thm. 1.13].

The intuitive idea described in the previous paragraph is placed into a rigorous framework by means of our central Assumption 2. In this context, it is worth noting that the actual main result, Theorem 2, couples the approximation accuracy t with the desired reconstruction error. This might appear somewhat peculiar at first sight because the processes $\{\xi_{t,i}(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ and $\{\varepsilon_{t,i}(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ usually become less regular with t decreasing. However, we demonstrate that this trade-off can be balanced out well and only affects the achievable oversampling rate.

Finally, our general technique to control empirical product processes of the form (1.3) may find application beyond signal recovery problems. For example, it could be also used to derive a new generation of nonlinear embedding results. In fact, one of the earliest works on binary random embeddings [50] has pointed out the usefulness of approximating nonlinearities, though strongly tailored to 1-bit quantization. With our systematic approach, their result could not only be reproduced but further improved (cf. [58, Thm. 3.7]).

1.4 Overview and Notation

The rest of this article is organized as follows. Section 2 presents our findings in full generality, where the proof of Theorem 2 is postponed to Sect. 7. In Sect. 3, we

³ This is an important difference to the non-uniform regime, where \mathbf{x} is a fixed signal. In this case, it is possible to apply known multiplier concentrations inequalities to control (1.3), e.g., see [42, Thm. 1.9].

elaborate aspects that are specifically due to uniform recovery and highlight potential gaps to the non-uniform regime. Sections 4 and 5 are then devoted to applications and concrete examples, including a series of new guarantees for nonlinear compressed sensing and high-dimensional estimation; the corresponding proofs are provided in Sects. 8 and 9, respectively. Our concluding discussion can be found in Sect. 6.

Before proceeding, let us fix some standard notations and conventions that are commonly used in this work. The letters c and C are reserved for (positive) constants, whose values could change from time to time. We speak of a *universal constant* if its value does not depend on any other involved parameter. If an inequality holds up to a universal constant C , we usually write $A \lesssim B$ instead of $A \leq C \cdot B$. The notation $A \asymp B$ is a shortcut for $A \lesssim B \lesssim A$.

For $p \in \mathbb{N}$, we set $[p] := \{1, \dots, p\}$. The *cardinality* of an index set $\mathcal{S} \subset [p]$ is denoted by $|\mathcal{S}|$ and its *set complement* in $[p]$ is $\mathcal{S}^c := [p] \setminus \mathcal{S}$. Vectors and matrices are denoted by lower- and uppercase boldface letters, respectively. The j -th entry of a vector $\mathbf{v} \in \mathbb{R}^p$ is denoted by $(\mathbf{v})_j$, or simply by v_j if there is no danger of confusion. The *support* of $\mathbf{v} \in \mathbb{R}^p$ is defined by $\text{supp}(\mathbf{v}) := \{j \in [p] : v_j \neq 0\}$ and $\|\mathbf{v}\|_0 := |\text{supp}(\mathbf{v})|$ denotes its *sparsity*. We write $\mathbf{I}_p \in \mathbb{R}^{p \times p}$ and $\mathbf{0} \in \mathbb{R}^p$ for the *identity matrix* and the *zero vector* in \mathbb{R}^p , respectively. For $1 \leq q \leq \infty$, we denote the ℓ^q -norm on \mathbb{R}^p by $\|\cdot\|_q$ and the associated *unit ball* by \mathbb{B}_q^p . The *Euclidean unit sphere* is given by $\mathbb{S}^{p-1} := \{\mathbf{v} \in \mathbb{R}^p : \|\mathbf{v}\|_2 = 1\}$. The spectral matrix norm is denoted by $\|\cdot\|_{2 \rightarrow 2}$.

Let $H, H' \subset \mathbb{R}^p$ and $\mathbf{v} \in \mathbb{R}^p$. We write $\text{span}(H)$ for the *linear hull* of H and the *linear cone* generated by H (not necessarily convex) is denoted by $\text{cone}(H) := \{s\tilde{\mathbf{v}} : \tilde{\mathbf{v}} \in H, s \geq 0\}$. By $P_H: \mathbb{R}^p \rightarrow \mathbb{R}^p$, we denote the *Euclidean projection* onto the closure of H if it is well defined, and we use the shortcut $P_{\mathbf{v}} := P_{\text{span}(\{\mathbf{v}\})}$. We write H^\perp for the *orthogonal complement* of H . The *Minkowski difference* between H and H' is defined by $H - H' := \{\mathbf{v}_1 - \mathbf{v}_2 : \mathbf{v}_1 \in H, \mathbf{v}_2 \in H'\}$, and we use the shortcut $H - \mathbf{v} := H - \{\mathbf{v}\}$. For $\varepsilon > 0$, we denote the *covering number* of H at scale ε with respect to the Euclidean norm by $\mathcal{N}(H, \varepsilon)$.

The L^q -norm of a real-valued random variable a is given by $\|a\|_{L^q} := (\mathbb{E}[|a|^q])^{1/q}$. We call a *sub-Gaussian* if

$$\|a\|_{\psi_2} := \inf \{v > 0 : \mathbb{E}[\exp(|a|^2/v^2)] \leq 2\} < \infty,$$

and $\|\cdot\|_{\psi_2}$ is called the *sub-Gaussian norm*. A random vector \mathbf{a} in \mathbb{R}^p is called *sub-Gaussian* if $\|\mathbf{a}\|_{\psi_2} := \sup_{\mathbf{v} \in \mathbb{S}^{p-1}} \|\langle \mathbf{a}, \mathbf{v} \rangle\|_{\psi_2} < \infty$. We say that \mathbf{a} is *centered* if $\mathbb{E}[\mathbf{a}] = \mathbf{0}$, and it is *isotropic* if $\mathbb{E}[\mathbf{a}\mathbf{a}^T] = \mathbf{I}_p$. We write $\mathbf{a} \sim \mathbf{N}(\mathbf{0}, \mathbf{I}_p)$ if \mathbf{a} is a *standard Gaussian random vector* in \mathbb{R}^p . For a more detailed introduction to sub-Gaussian random variables and their properties, see [64, Chap. 2 & 3]. Now, let (\mathcal{H}, d) be a pseudo-metric space and consider a real-valued stochastic process $\{a_h\}_{h \in \mathcal{H}}$ on \mathcal{H} . Then, $\{a_h\}_{h \in \mathcal{H}}$ has *sub-Gaussian increments* with respect to d if

$$\|a_h - a_{h'}\|_{\psi_2} \leq d(h, h') \quad \text{for all } h, h' \in \mathcal{H}. \tag{1.4}$$

We say that a function $f: \mathbb{R}^p \rightarrow \mathbb{R}^{p'}$ is γ -Lipschitz if it is Lipschitz continuous with respect to the Euclidean metric and a Lipschitz constant $\gamma \geq 0$. The *sign* of $v \in \mathbb{R}$ is denoted by $\text{sign}(v)$, with the convention that $\text{sign}(0) := 1$. If $\text{sign}(\cdot)$ is applied to a

vector, the operation is understood entrywise. The *ceiling* and *floor function* of $v \in \mathbb{R}$ is denoted by $\lceil v \rceil$ and $\lfloor v \rfloor$, respectively.

2 Main Result

As preliminary steps, we introduce the formal model setup in Sect. 2.1, followed by our increment conditions on the observation variable in Sect. 2.2. The main result of Theorem 2 is then stated and discussed in Sect. 2.3, while its proof can be found in Sect. 7.

2.1 Model Setup

The following model assumption fixes the notation for the remainder of this section and specifies the instance of Problem 1 that we will study from now on:

Assumption 1 (a) Components of the observation model:

- *Measurement vector*: a centered, isotropic, sub-Gaussian random vector $\mathbf{a} \in \mathbb{R}^p$ such that $\|\mathbf{a}\|_{\psi_2} \leq L$ for some $L > 0$.⁴
- *Signal set*: a set \mathcal{X} (not necessarily a subset of \mathbb{R}^p).
- *Output function*: a scalar function $F : \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ that may be random (not necessarily independent of \mathbf{a}).
- *Observation variable*: $\tilde{y}(\mathbf{x}) := F(\mathbf{a}, \mathbf{x})$ for $\mathbf{x} \in \mathcal{X}$.

(b) Components of the measurement process:

- *Measurement ensemble*: $\{(\mathbf{a}_i, F_i)\}_{i=1}^m$ are independent copies of the *measurement model* (\mathbf{a}, F) .
- *Observations*: $\tilde{y}_i(\mathbf{x}) := F_i(\mathbf{a}_i, \mathbf{x})$ for $i = 1, \dots, m$ and $\mathbf{x} \in \mathcal{X}$. The *observation vector* of \mathbf{x} is denoted by $\tilde{\mathbf{y}}(\mathbf{x}) := (\tilde{y}_1(\mathbf{x}), \dots, \tilde{y}_m(\mathbf{x})) \in \mathbb{R}^m$.

(c) Components of the recovery procedure:

- *Constraint set*: a convex subset $K \subset \mathbb{R}^p$ (not necessarily bounded).
- *Target function*: a map $T : \mathcal{X} \rightarrow K$, such that $T\mathcal{X} \subset K$ is bounded.

Part (a) of Assumption 1 establishes the statistical “template” of our measurement model, from which i.i.d. observations are drawn according to part (b). Importantly, this defines an entire class of observation variables $\{\tilde{y}(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$, which corresponds to a real-valued stochastic process. Assumption 1(c) is arguably the most abstract part, but forms a crucial ingredient of our main result: Theorem 2 states an error bound for $T\mathbf{x} \in K$ instead of the actual signal $\mathbf{x} \in \mathcal{X}$ (which might not even belong to \mathbb{R}^p). The basic linearization strategy behind $(\mathbf{P}_{K, \mathbf{y}})$ typically calls for an appropriate transformation of the signal domain. A good example is the special case of Theorem 1, where T amounts to a rescaling of \mathbf{x} by the model-dependent factor μ . Such a simple correction is sufficient for most applications considered in Sect. 4; but there exist more complicated situations, such as the variable selection problem in Sect. 4.5, where \mathcal{X}

⁴ Note that we actually have that $L \geq \sqrt{1/\log(2)} > 1$ if \mathbf{a} is isotropic, see [35].

represents support sets on $[p]$. At the present level of abstraction, it is useful to think of $T\mathbf{x}$ as a parameterized version of $\mathbf{x} \in \mathcal{X}$ that is compatible with the estimation procedure of $(\mathbf{P}_{K,\mathbf{y}})$ and its constraint set $K \subset \mathbb{R}^p$.⁵

In principle, T can be arbitrary in Assumption 1, but there often exists a canonical choice that is driven by the size of the following model parameter:

Definition 1 Let Assumption 1 be satisfied. Then, we define the *target mismatch* of $\mathbf{x} \in \mathcal{X}$ by

$$\rho(\mathbf{x}) := \|\mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] - T\mathbf{x}\|_2.$$

The meaning of this definition becomes clearer when taking the perspective of statistical learning for a moment. Assuming that $y_i = \tilde{y}_i(\mathbf{x})$, the Lasso estimator $(\mathbf{P}_{K,\mathbf{y}})$ can be viewed as an empirical loss minimization problem with “data” $\{(\mathbf{a}_i, y_i)\}_{i=1}^m$. A central question is then under what conditions its solution can approximate (in a proper sense) the associated expected loss minimization problem (obtained in the infinite sample limit $m \rightarrow \infty$):

$$\min_{z \in K} \mathbb{E}[(\tilde{\mathbf{y}}(\mathbf{x}) - \langle \mathbf{a}, z \rangle)^2]. \tag{2.1}$$

It is not hard to see that the only critical point of the objective function in (2.1) is the vector $\mathbf{x}^* := \mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}]$, and if $\mathbf{x}^* \in K$, this is the global optimum. Therefore, and as its name suggests, $\rho(\mathbf{x})$ measures the mismatch between the target vector $T\mathbf{x}$ and the (global) expected loss minimizer \mathbf{x}^* .

In the context of Theorem 2, the target mismatch can be seen as an upper bound for the asymptotic error of estimating $T\mathbf{x}$ via $(\mathbf{P}_{K,\mathbf{y}})$. Consequently, a general rule of thumb is to select $T\mathbf{x}$ such that $\rho(\mathbf{x})$ vanishes or becomes sufficiently small. However, simply setting $T\mathbf{x} := \mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}]$ (if contained in K) is not necessarily consistent with our wish for signal recovery, e.g., in the application in Sect. 4.2. Therefore, we have left T unspecified in Assumption 1, even though a proper choice will be obvious for all considered examples.

We close this subsection with a short remark about the noise model considered in this work:

Remark 1 (Noise model) We will adopt the *adversarial* (or *worst-case*) noise model, which is the common setting in the field of compressed sensing, cf. [21]. More specifically, when interested in recovery of $\mathbf{x} \in \mathcal{X}$, the actual input $\mathbf{y} = (y_1, \dots, y_m)$ for $(\mathbf{P}_{K,\mathbf{y}})$ need not exactly correspond to the observation vector $\tilde{\mathbf{y}}(\mathbf{x}) = (\tilde{y}_1(\mathbf{x}), \dots, \tilde{y}_m(\mathbf{x}))$ introduced in Assumption 1(b).⁶ Instead, arbitrary perturbations are allowed as long as $d(\mathbf{y}, \tilde{\mathbf{y}}(\mathbf{x})) \lesssim t$ for an appropriate noise metric $d(\cdot, \cdot)$ and a desired accuracy t . In Theorem 1, for instance, we have $\tilde{y}_i(\mathbf{x}) = f(\langle \mathbf{a}_i, \mathbf{x} \rangle)$ and

⁵ One should bear in mind that the target function T is a purely *theoretical* object that does not affect the solution of $(\mathbf{P}_{K,\mathbf{y}})$ and may be (partially) unknown in practice. Nevertheless, it is an essential analysis tool for Problem 1, relating the underlying observation model to the actual recovery method $(\mathbf{P}_{K,\mathbf{y}})$, see also Remark 2(1).

⁶ The use of the modifier ‘ \sim ’ in Assumption 1 is also due to this fact and may help to distinguish between our observation model and the actual input for $(\mathbf{P}_{K,\mathbf{y}})$.

$y_i = \tilde{y}_i(\mathbf{x}) + v_i$, while $d(\cdot, \cdot)$ corresponds to the normalized ℓ^2 -norm. Thus, our overall goal is to show that, given $\{(\mathbf{a}_i, F_i)\}_{i=1}^m$, recovery is possible for *any* $\mathbf{x} \in \mathcal{X}$ and *any* moderate perturbation of its observation $\tilde{\mathbf{y}}(\mathbf{x})$.

In this context, it is worth noting that Assumption 1(b) does not require the variables $\tilde{y}_i(\mathbf{x})$ to be deterministic when conditioned on \mathbf{a}_i , since F_i itself may be random. For example, one could model *additive random noise* in that way, as common in statistical estimation theory. In principle, such a type of noise is also admissible in our main result, Theorem 2, but it is not entirely compatible with the aforementioned goal of *uniform* recovery. Indeed, the ensemble $\{(\mathbf{a}_i, F_i)\}_{i=1}^m$ is only drawn once, so that a single noise realization has to be considered for all $\mathbf{x} \in \mathcal{X}$. We therefore mostly stick to the adversarial perspective; see Remark 4(2) in Sect. 3 for a more detailed comparison of both noise models in the situation of linear measurements. Finally, we point out that randomized output functions are nevertheless very useful in the uniform regime, such as in the design of dithering variables (see Sect. 4.2).

2.2 Analysis Strategy and Increment Conditions

From now on, let Assumption 1 be satisfied. In this subsection, we introduce the key condition for our uniform recovery guarantee, namely that the class of observation variables $\{\tilde{y}(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$, or at least an approximation thereof, has sub-Gaussian increments. To better understand the relevance of this condition, it is insightful to first take a closer look at our basic analysis strategy for the generalized Lasso:

$$\min_{\mathbf{z} \in K} \frac{1}{m} \sum_{i=1}^m (y_i - \langle \mathbf{a}_i, \mathbf{z} \rangle)^2. \tag{P_{K,y}}$$

Note that we hide the dependency on the measurement vectors $\{\mathbf{a}_i\}_{i=1}^m$ when referring to $(P_{K,y})$, as this will be always clear from the context. At the present stage, the vector $\mathbf{y} = (y_1, \dots, y_m) \in \mathbb{R}^m$ in $(P_{K,y})$ is unspecified, but taking the viewpoint of our main result, Theorem 2, is already useful: If $\mathbf{x} \in \mathcal{X}$ is to be recovered, then \mathbf{y} is a noisy version of the observation vector $\tilde{\mathbf{y}}(\mathbf{x}) = (\tilde{y}_1(\mathbf{x}), \dots, \tilde{y}_m(\mathbf{x}))$ as defined in Assumption 1(b); see also Remark 1.

We now derive a simple, yet important criterion for an error bound for a (fixed) signal $\mathbf{x} \in \mathcal{X}$. To this end, let $\bar{\mathcal{L}}_x(\mathbf{v}) := \frac{1}{m} \sum_{i=1}^m (y_i - \langle \mathbf{a}_i, \mathbf{v} + T\mathbf{x} \rangle)^2$ denote the *empirical loss* of $\mathbf{v} \in \mathbb{R}^p$ over \mathbf{x} . For the sake of notational convenience, we have fixed the anchor point $T\mathbf{x}$ here, so that $(P_{K,y})$ is equivalent to minimizing $\bar{\mathcal{L}}_x(\cdot)$ over $K - T\mathbf{x}$. The *excess loss* of $\mathbf{v} \in \mathbb{R}^p$ over \mathbf{x} is then defined by

$$\mathcal{E}_x(\mathbf{v}) := \bar{\mathcal{L}}_x(\mathbf{v}) - \bar{\mathcal{L}}_x(\mathbf{0}).$$

It measures how much the empirical loss is changing when traveling from $T\mathbf{x}$ in the direction of \mathbf{v} . The following fact is an immediate consequence of the convexity of K and $\bar{\mathcal{L}}_x$; see Fig. 1 for an illustration.




Fig. 1 Illustration of the localization argument behind Fact 1. We assume that $\mathcal{E}_x(\mathbf{v}) > 0$ for all $\mathbf{v} \in K_{x,t}$, or equivalently, $\mathcal{E}_x(\mathbf{z} - T\mathbf{x}) > 0$ for all $\mathbf{z} \in K_{x,t} + T\mathbf{x}$ (see red arc). Since $\mathcal{E}_x(\mathbf{0}) = 0$, the convexity of K and $\mathcal{E}_x(\cdot)$ implies that a minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ must not lie outside the dark gray intersection. In other words, we have that $\|\hat{\mathbf{z}} - T\mathbf{x}\|_2 \leq t$

Fact 1 For $\mathbf{x} \in \mathcal{X}$ and a desired reconstruction accuracy $t > 0$, we set

$$K_{x,t} := \{\mathbf{v} \in K - T\mathbf{x} : \|\mathbf{v}\|_2 = t\} = (K - T\mathbf{x}) \cap t\mathbb{S}^{p-1}.$$

If $\mathcal{E}_x(\mathbf{v}) > 0$ for all $\mathbf{v} \in K_{x,t}$, then every minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - T\mathbf{x}\|_2 \leq t$.

For a fixed accuracy $t > 0$, Fact 1 implies uniform recovery for all those $\mathbf{x} \in \mathcal{X}$ satisfying $\inf_{\mathbf{v} \in K_{x,t}} \mathcal{E}_x(\mathbf{v}) > 0$. This motivates us to establish a uniform lower bound for the excess loss for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{v} \in K_{x,t}$. In particular, such a task is considerably more difficult than showing a bound on $K_{x,t}$ for a fixed $\mathbf{x} \in \mathcal{X}$, which would be in accordance with the non-uniform results discussed in the second part of Sect. 1.1. To get a better sense of this challenge, let us consider the following basic decomposition of the excess loss:

$$\begin{aligned} \mathcal{E}_x(\mathbf{v}) &= \underbrace{\frac{1}{m} \sum_{i=1}^m \langle \mathbf{a}_i, \mathbf{v} \rangle^2}_{=: \mathcal{Q}(\mathbf{v})} - \underbrace{\frac{2}{m} \sum_{i=1}^m (y_i - \tilde{y}_i(\mathbf{x})) \langle \mathbf{a}_i, \mathbf{v} \rangle}_{=: \mathcal{N}_x(\mathbf{v})} \\ &\quad + \underbrace{\frac{2}{m} \sum_{i=1}^m (\langle \mathbf{a}_i, T\mathbf{x} \rangle - \tilde{y}_i(\mathbf{x})) \langle \mathbf{a}_i, \mathbf{v} \rangle}_{=: \mathcal{M}_x(\mathbf{v})}. \end{aligned} \tag{2.2}$$

The quadratic term $\mathcal{Q}(\mathbf{v})$ and noise term $\mathcal{N}_x(\mathbf{v})$ are rather unproblematic and can be controlled by a recent matrix deviation inequality (see Step 1 and Step 2 in Sect. 7). Much more intricate is the multiplier term $\mathcal{M}_x(\mathbf{v})$, since the underlying multiplier variable $\xi(\mathbf{x}) := \langle \mathbf{a}, T\mathbf{x} \rangle - \tilde{y}(\mathbf{x})$ depends on \mathbf{x} , so that we actually have to deal with an empirical product process. A major difficulty is that known concentration results for such product processes do not apply directly, since the class $\{\xi(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ need

not have sub-Gaussian increments with respect to an appropriate (pseudo-)metric. For example, this would happen if we would drop the Lipschitz assumption on f in Theorem 1. The key idea of our approach is to approximate $\tilde{y}(\mathbf{x})$ by a more “regular” observation variable $\tilde{y}_t(\mathbf{x})$ in such a way that the resulting multiplier variable and the approximation error both have sub-Gaussian increments. This is made precise by the following assumption:

Assumption 2 We define a pseudo-metric on \mathcal{X} by $d_T(\mathbf{x}, \mathbf{x}') := \|T\mathbf{x} - T\mathbf{x}'\|_2$ for $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$. For $t \geq 0$, assume that there exists a class of observation variables $\{\tilde{y}_t(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ such that the following properties hold:

- (a) *Approximation error:* Setting $\varepsilon_t(\mathbf{x}) := |\tilde{y}(\mathbf{x}) - \tilde{y}_t(\mathbf{x})|$, we assume that $\mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |(\mathbf{a}, \mathbf{z})|] \leq \frac{t}{64}$ for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathbb{S}^{p-1}$.
- (b) *Multiplier increments:* Set $\xi_t(\mathbf{x}) := (\mathbf{a}, T\mathbf{x}) - \tilde{y}_t(\mathbf{x})$ for $\mathbf{x} \in \mathcal{X}$. We assume that there exist $r \geq 0$ and $L_t \geq 0$ such that

$$\|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} \leq L_t \cdot d_T(\mathbf{x}, \mathbf{x}') \quad \text{and} \quad \|\xi_t(\mathbf{x})\|_{\psi_2} \leq r \quad \text{for all } \mathbf{x}, \mathbf{x}' \in \mathcal{X}.$$

- (c) *Error increments:* We assume that there exist $\hat{r} \geq 0$ and $\hat{L}_t \geq 0$ such that

$$\|\varepsilon_t(\mathbf{x}) - \varepsilon_t(\mathbf{x}')\|_{\psi_2} \leq \hat{L}_t \cdot d_T(\mathbf{x}, \mathbf{x}') \quad \text{and} \quad \|\varepsilon_t(\mathbf{x})\|_{\psi_2} \leq \hat{r} \quad \text{for all } \mathbf{x}, \mathbf{x}' \in \mathcal{X}.$$

Assumption 2(b) and (c) imply that $\{\xi_t(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ and $\{\varepsilon_t(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ have sub-Gaussian increments with respect to $L_t \cdot d_T$ and $\hat{L}_t \cdot d_T$, respectively. Similarly, r and \hat{r} bound the (sub-Gaussian) diameter of the respective classes. A convenient interpretation of Assumption 2 is as follows: Choose an approximation $\tilde{y}_t(\mathbf{x})$ for every $\mathbf{x} \in \mathcal{X}$ such that the error does not become too large in the sense of part (a); at the same time, ensure that the increment conditions of part (b) and (c) are satisfied such that the parameters L_t and \hat{L}_t do not grow too fast as t becomes smaller. A remarkable conclusion from our applications to quantized compressed sensing in Sects. 4.1–4.3 is that this strategy can even succeed for observation variables with discontinuous output functions.

Finally, we emphasize that the approximation error $\varepsilon_t(\mathbf{x})$ includes taking the absolute value, which is crucial to our proof (see Step 3 in Sect. 7). In particular, Assumption 2 does not necessarily imply that the original class $\{\xi(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ has sub-Gaussian increments with respect to $\tilde{L} \cdot d_T$ for some $\tilde{L} \geq 0$. On the other hand, it is certainly possible that $\{\xi(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ already has sub-Gaussian increments, such as in Theorem 1. In this case, we can simply choose $\tilde{y}_t(\mathbf{x}) := \tilde{y}(\mathbf{x})$, so that $\varepsilon_t(\mathbf{x}) = 0$ and Assumption 2(a) and (c) are trivially fulfilled.

2.3 Uniform Recovery Guarantee

In order to formulate the main result of this work, we require the notion of Gaussian mean width. This geometric parameter has proven to be a useful complexity measure in high-dimensional signal recovery, e.g., see [2, 12, 43, 54, 57] for pioneering works in that direction. Our approach is no exception and makes use of the following localized version.

Definition 2 Let $H \subset \mathbb{R}^p$ and $\mathbf{g} \sim \mathbf{N}(\mathbf{0}, \mathbf{I}_p)$. The (Gaussian) mean width of H is given by

$$w(H) := \mathbb{E} \left[\sup_{\mathbf{v} \in H} \langle \mathbf{g}, \mathbf{v} \rangle \right].$$

For $t \geq 0$, we define the local mean width of H (at scale t) by

$$w_t(H) := \begin{cases} w(\frac{1}{t}H \cap \mathbb{S}^{p-1}), & t > 0, \\ w(\text{cone}(H) \cap \mathbb{S}^{p-1}), & t = 0. \end{cases}$$

As final preparatory step, we introduce two expressions that capture the size of measurement noise: For $m_0 \in \{0, 1, \dots, m\}$ and $\mathbf{w} = (w_1, \dots, w_m) \in \mathbb{R}^m$, let

$$\|\mathbf{w}\|_{[m_0]} := \left(\sum_{i=1}^{m_0} |w_i^*|^2 \right)^{1/2} \quad \text{and} \quad \sigma_{m_0}(\mathbf{w})_2 := \left(\sum_{i=m_0+1}^m |w_i^*|^2 \right)^{1/2},$$

where (w_1^*, \dots, w_m^*) is the non-increasing rearrangement of $(|w_1|, \dots, |w_m|)$. Obviously, we have that $\|\mathbf{w}\|_{[m_0]}^2 + \sigma_{m_0}(\mathbf{w})_2^2 = \|\mathbf{w}\|_2^2$, and in particular, $\|\mathbf{w}\|_{[0]} = 0$ and $\sigma_0(\mathbf{w})_2 = \|\mathbf{w}\|_2$. Moreover, note that $\|\cdot\|_{[m_0]}$ simply corresponds to the ℓ^2 -norm of the m_0 -largest entries, while $\sigma_{m_0}(\cdot)_2$ is commonly known as the ℓ^2 -error of the best m_0 -term approximation.

We are now ready to state our main recovery guarantee, which forms the basis of all applications presented in Sect. 4; see Sect. 7 for a complete proof.

Theorem 2 *There exist universal constants $c, C > 0$ for which the following holds.*

Let Assumptions 1 and 2 be satisfied for a fixed accuracy $t \geq 0$, and let $m_0 \in \{0, 1, \dots, m\}$. For $\Delta > 0, u \geq 1$, and $u_0 \geq \sqrt{m_0 \log(em/m_0)}$, we assume that^{7 8}

$$m \geq C \cdot L^2 \cdot \max \left\{ \log L \cdot (w_t^2(K - T\mathcal{X}) + u^2), \right. \\ \left. L^2 \Delta^2 \cdot (w_t^2(K - T\mathcal{X}) + u^2), \right. \\ \left. t^{-2}(r^2 + \hat{r}^2) \cdot (w_t^2(K - T\mathcal{X}) + u^2) + t^{-2}(L_t^2 + \hat{L}_t^2) \cdot w^2(T\mathcal{X}) \right\}. \quad (2.3)$$

Then with probability at least $1 - \exp(-cu^2) - \exp(-cu_0^2)$ on the random draw of $\{(\mathbf{a}_i, F_i)\}_{i=1}^m$, the following holds uniformly for every $\mathbf{x} \in \mathcal{X}$ with $\rho(\mathbf{x}) \leq \frac{t}{32}$: Let $\mathbf{y} \in \mathbb{R}^m$ be any input vector such that

$$\frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_{[m_0]} \leq \Delta t \quad \text{and} \quad \frac{1}{\sqrt{m}} \sigma_{m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}))_2 \leq \frac{t}{20}. \quad (2.4)$$

Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - T\mathbf{x}\|_2 \leq t$.

⁷ The condition (2.3) is stated in such a way that it is convenient to handle in our specific applications (see Sect. 4). However, for a basic understanding, the reader may simply set $u = u_0$, so that the first and second branch in (2.3) can be merged to $m \gtrsim L^2(\log L + L^2 \Delta^2) \cdot (w_t^2(K - T\mathcal{X}) + u^2)$.

⁸ In the case of exact recovery, i.e., $t = 0$, we follow the convention $0 \cdot \infty := 0$, so that the condition (2.3) requires that $r = \hat{r} = L_t = \hat{L}_t = 0$. Then, (2.3) is particularly fulfilled for $m \geq C \cdot L^2 \log L \cdot (w_0^2(K - T\mathcal{X}) + u^2)$, where $m_0 = 0, \Delta = L^{-1} \sqrt{\log L}$, and $u = u_0$.

The constraints of (2.4) imply that the estimator $(P_{K,y})$ is *robust* against (adversarial) perturbations and outliers in the observation vector $\tilde{y}(x)$. Here, a central role is played by the fine-tuning parameter m_0 , controlling which coefficients of the noise vector $y - \tilde{y}(x)$ are captured by the first and second condition in (2.4), respectively. Enlarging m_0 helps to suppress “gross” outliers (that may appear in the m_0 largest coefficients of $y - \tilde{y}(x)$ in magnitude). Indeed, by adjusting the free parameter Δ , the first condition of (2.4) becomes less restrictive than the second one (which captures the remaining coefficients of $y - \tilde{y}(x)$). On the other hand, choosing m_0 and/or Δ too large may have a negative effect on the sample complexity in the second branch of (2.3). The benefit of permitting outliers will become clearer in the context of quantization where the noise is due to wrong bits in y instead of real-valued perturbations (see Sects. 4.1–4.3). Note that the most common noise model is obtained for $m_0 = 0$: (2.4) then simply corresponds to the baseline condition $\frac{1}{\sqrt{m}} \|y - \tilde{y}(x)\|_2 \leq \frac{t}{20}$, which is consistent with standard noise bounds from the literature. We refer to Remark 4(2) in the next section for a more detailed comparison of adversarial and statistical noise, as well as aspects of optimality.

The second important constraint of Theorem 2 is that it does not apply to those $x \in \mathcal{X}$ with $\rho(x) > \frac{t}{32}$. Hence, the target mismatch $\rho(x)$ can be seen as an upper bound for the asymptotic error when estimating Tx via $(P_{K,y})$. In particular, the above error bound does not allow for arbitrarily high precision unless $\rho(x) = 0$.

Let us now turn to the key assumption (2.3), which relates the number of measurements m to the desired accuracy t . The above formulation views m as a function of t , specifying how fast m grows if t decreases. The inverse relationship is also of interest, since it describes the reconstruction error as a function of m . However, this can only be made explicit in specific situations where the parameters at the right-hand side of (2.3) are known or can be well estimated. Of special importance in that respect are the increment parameters L_t and \hat{L}_t from Assumption 2, whose size has a significant impact on the required (over-)sampling rate. For 1-bit observations as studied in Sect. 4.1, for example, we achieve $L_t, \hat{L}_t \lesssim t^{-1}$, so that (2.3) would yield an error decay rate of $O(m^{-1/4})$ —or conversely, an oversampling rate of $O(t^{-4})$. Remarkably, we will derive a corollary of Theorem 2 in Sect. 5 that can bypass the restrictions of Assumption 2 at the price of a stronger (local) stability condition on the observation model. This allows us to combine our approach with known embedding results from the literature and thereby to obtain near-optimal decay rates in special cases.

The dependence of (2.3) on the complexity parameters $w_t(K - T\mathcal{X})$ and $w(T\mathcal{X})$ is quite natural. Similar expressions have already appeared in several of the articles discussed in Sect. 1.1, e.g., see [17, 49, 52]. A detailed discussion, including concrete examples and possible simplifications, can be found in Sect. 3. A distinctive feature of Theorem 2 is that the considered version of local mean width $w_t(K - T\mathcal{X})$ “compares” the parameter vectors in K only to those in the transformed signal set $T\mathcal{X} \subset K$. This refinement of the common quantity $w_t(K - K)$ can lead to improved guarantees in certain scenarios, namely when $T\mathcal{X}$ is much smaller than K (see Example 1 for an illustration of this aspect).

Remark 2 (1) *Inversion of T .* In principle, the target function T neither has to be injective nor does it have to be explicitly known to solve $(P_{K,y})$. However, in order

to obtain a practicable statement from Theorem 2, the actual signal $\mathbf{x} \in \mathcal{X}$ should be extractable from an approximation $\hat{\mathbf{z}}$ of $T\mathbf{x}$. An appropriate implementation of T can differ considerably from situation to situation. In the case of single-index models (see Theorem 1), for example, this would involve a rescaling of $\hat{\mathbf{z}}$, which requires a (rough) knowledge of the nonlinearity f . In contrast, for variable selection (see Sect. 4.5), it can be sufficient to perform a simple hard-thresholding step on $\hat{\mathbf{z}}$ to obtain a good estimate of the underlying support.

- (2) *Optimality.* The best possible error decay rate that can result from Theorem 2 is $O(m^{-1/2})$, supposed that we have $L_t, \hat{L}_t \lesssim 1$.⁹ This rate cannot be improved in general, or in other words, the exponent $-\frac{1}{2}$ is not an artifact of our proof but corresponds to a fundamental statistical barrier (see [52, Sec. 4]). However, it is possible to break through this barrier in specific model setups, e.g., *noiseless* 1-bit observations [34, 36]. Such superior rates are usually not achieved by $(P_{K,y})$, but may require a more sophisticated estimator instead.

Remark 3 (Possible extensions) For the sake of clarity, we did not present the most general version of Theorem 2 that we could have derived. There are several variants and generalizations that we expect to be practicable:

- (1) *Anisotropic measurements.* Based on an observation from [51, Rmk. 1.7], [22, Sec. II.D], the isotropy of the measurement vectors in Assumption 1(a) can be relaxed to any positive definite covariance matrix $\Sigma \in \mathbb{R}^{p \times p}$. To see this, let us assume that \mathbf{a} is centered and sub-Gaussian with $\mathbb{E}[\mathbf{a}\mathbf{a}^T] = \Sigma$. Then, we may write $\mathbf{a} = \sqrt{\Sigma}\bar{\mathbf{a}}$ where $\bar{\mathbf{a}}$ is isotropic, centered, and sub-Gaussian. If $\bar{\mathbf{a}}_1, \dots, \bar{\mathbf{a}}_m$ are i.i.d. copies of $\bar{\mathbf{a}}$, a simple reformulation of $(P_{K,y})$ yields

$$\operatorname{argmin}_{\mathbf{z} \in K} \frac{1}{m} \sum_{i=1}^m (y_i - \langle \mathbf{a}_i, \mathbf{z} \rangle)^2 = \sqrt{\Sigma}^{-1} \cdot \operatorname{argmin}_{\bar{\mathbf{z}} \in \sqrt{\Sigma}K} \frac{1}{m} \sum_{i=1}^m (y_i - \langle \bar{\mathbf{a}}_i, \bar{\mathbf{z}} \rangle)^2.$$

In other words, $(P_{K,y})$ is equivalent to a modified estimator that operates with an isotropic measurement ensemble. Therefore, we may apply Theorem 2 directly to the latter one, when replacing K and $T(\cdot)$ by $\sqrt{\Sigma}K$ and $\sqrt{\Sigma}T(\cdot)$ in Assumption 1, respectively. Note that these adaptations are of purely theoretical nature and in practice, no explicit knowledge of Σ is required. In particular, this procedure still implies an error bound for minimizers $\hat{\mathbf{z}}$ of $(P_{K,y})$:

$$\|\hat{\mathbf{z}} - T\mathbf{x}\|_2 \leq \|\sqrt{\Sigma}^{-1}\|_{2 \rightarrow 2} \cdot \|\sqrt{\Sigma}\hat{\mathbf{z}} - \sqrt{\Sigma}T\mathbf{x}\|_2 \leq \|\sqrt{\Sigma}^{-1}\|_{2 \rightarrow 2} \cdot t.$$

Similarly, one can conveniently control the adapted complexity parameters (see [51, Rmk. 1.7], [22, Sec. II.D]):

$$w_t(\sqrt{\Sigma}K - \sqrt{\Sigma}T\mathcal{X}) \leq \max\{1, \|\sqrt{\Sigma}^{-1}\|_{2 \rightarrow 2}\} \cdot \|\sqrt{\Sigma}\|_{2 \rightarrow 2} \cdot w(\frac{1}{t}(K - T\mathcal{X}) \cap \mathbb{B}_2^p),$$

⁹ Note that the expression $O(m^{-1/2})$ suppresses the dependence on $w_t(K - T\mathcal{X})$. Although the latter can be trivially bounded by $w_0(K - T\mathcal{X})$, which is independent of t , such an estimate might not appropriately capture the (low) complexity of K in certain cases.

$$w(\sqrt{\Sigma}T\mathcal{X}) \lesssim \|\sqrt{\Sigma}\|_{2 \rightarrow 2} \cdot w(T\mathcal{X}).$$

Overall, we can conclude that accurate (uniform) recovery is also possible with anisotropic measurements, as long as the (unknown) covariance matrix is well conditioned.

- (2) *Even output functions.* There exist certain types of nonlinear models that are not directly compatible with the estimator $(P_{K,y})$. For example, if $f : \mathbb{R} \rightarrow \mathbb{R}$ is an even function in Theorem 1, it turns out that $\mu = \mathbb{E}_{g \sim N(0,1)}[f(g)g] = 0$. In other words, $(P_{K,y})$ would simply approximate the zero vector, which is clearly not what one is aiming at. Notably, this scenario includes the important phase-retrieval problem $f(\cdot) = |\cdot|$. Fortunately, the linearization idea behind $(P_{K,y})$ is still applicable in these situations, when combined with the *phase-lifting* trick [9]:

$$\min_{X \in K} \frac{1}{m} \sum_{i=1}^m (y_i - \langle a_i a_i^T - \mathbb{E}[a_i a_i^T], X \rangle_F)^2, \tag{P_{K,y}^{(lifted)}}$$

where $\langle \cdot, \cdot \rangle_F$ denotes the Hilbert–Schmidt inner product and $K \supset \{xx^T : x \in \mathcal{X}\}$ is a convex subset of the positive semi-definite cone in $\mathbb{R}^{p \times p}$. The recovery performance of $(P_{K,y}^{(lifted)})$ for even output functions has been recently analyzed by the first author in [25, Sec. 3.5], inspired by earlier findings of Thrampoulidis and Rawat [61]. A key challenge in this regard is that the lifted measurement vectors/matrices $a_i a_i^T - \mathbb{E}[a_i a_i^T]$ are *heavier tailed* than a_i , which in turn has required deeper insights from generic chaining theory. We expect that our results are extendable into this direction, but with a technical effort that would go beyond the scope of this work.

- (3) *Different metrics.* An extension that is specific to uniform recovery concerns the choice of pseudo-metric in Assumption 2. In principle, d_T could be replaced by an arbitrary pseudo-metric d on the signal set \mathcal{X} . The claim and proof of Theorem 2 would still hold true if the (global) complexity of \mathcal{X} is captured by $\gamma_2(\mathcal{X}, d)$ instead of $w(T\mathcal{X})$; see Step 3 in Sect. 7. While this would make our approach slightly more flexible, we note that Talagrand’s γ_2 -functional is difficult to control in general. Apart from that, the ℓ^2 -norm as *error metric* in Theorem 2 could be replaced by an arbitrary semi-norm $\|\cdot\|$. Specifically, this step would lead to an error bound in terms of $\|\cdot\|$ in Fact 1, but it also entails an adaption of the spherical intersection in the local mean width (see Definition 2). Finally, different *convex loss functions* are feasible for $(P_{K,y})$ as well [22], although such an extension would require several additional technical assumptions.

3 (Non-)Uniformity and Signal Complexity

In this section, we highlight important aspects that are specifically attributable to uniform recovery, compared to existing non-uniform results. This includes a more detailed discussion of the local mean width $w_t(K - T\mathcal{X})$, which appears as the central complexity parameter in our approach.

Except for the inclusion $T\mathcal{X} \subset K$, Theorem 2 does not impose any restrictions on the signal set \mathcal{X} . If T is the identity function, the choice of \mathcal{X} may range from a singleton $\mathcal{X} = \{\mathbf{x}\}$ to the entire constraint set $\mathcal{X} = K$. The latter scenario corresponds to uniform recovery of all signals in K , while the former means that one is interested in a specific $\mathbf{x} \in K$ —which is nothing else than non-uniform recovery. In particular, due to $w(\{\mathbf{x}\}) = 0$, Theorem 2 is consistent with previously known non-uniform guarantees, e.g., see [23, Thm. 3.6] and [51, Thm 1.9]. Our main result therefore also indicates what additional expenses may come in the case of uniform recovery.

Let us first focus on the role of *nonlinear* observations, whose impact is reflected by the third branch of (2.3) in Theorem 2. The key quantities in this respect are the increment parameters L_t and \hat{L}_t . Their size strongly depends on the type of output function F considered in Assumption 1. For instance, a jump discontinuity in F can cause a weaker oversampling rate, in the sense that $L_t, \hat{L}_t \lesssim t^{-1}$; we refer to Sects. 4.1–4.4 for concrete examples and Sect. 5 for a possible improvement. In the non-uniform case ($\mathcal{X} = \{\mathbf{x}\}$), on the other hand, this issue becomes irrelevant because one can simply set $L_t = \hat{L}_t = 0$ (since Assumption 2(b) and (c) need only be satisfied for a single point). Thus, we can draw the following informal conclusion:

The transition to uniform recovery with nonlinear output functions may result in a worse oversampling rate (with respect to t).

The second foundational aspect of uniformity concerns the geometric complexity of the constraint set K and its interplay with the actual signal set \mathcal{X} . To keep our exposition as clear as possible, we restrict ourselves to the situation of linear observations (though most arguments naturally carry over to nonlinear models). Applying Theorem 2 to this special case leads to the following recovery guarantee (see Sect. 10 for a proof). Note that we intentionally allow for two different types of noise here, based on the adversarial ($\sim v_i$) and statistical model ($\sim \tau_i$); see Remark 4(2) for further discussion.

Corollary 1 *There exist universal constants $c, C > 0$ for which the following holds.*

Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent copies of a centered, isotropic, sub-Gaussian random vector $\mathbf{a} \in \mathbb{R}^p$ with $\|\mathbf{a}\|_{\psi_2} \leq L$. Let τ_1, \dots, τ_m be independent copies of $\tau \sim N(0, \sigma^2)$ for some $\sigma \geq 0$ (also independent of $\{\mathbf{a}_i\}_{i=1}^m$). Moreover, let $\mathcal{X} \subset \mathbb{R}^p$ be a bounded subset and let $K \subset \mathbb{R}^p$ be a convex set such that $\mathcal{X} \subset K$. For $u \geq 1$ and $t \geq 0$, we assume that

$$m \geq C \cdot L^2 \cdot (\log L + t^{-2}\sigma^2) \cdot (w_t^2(K - \mathcal{X}) + u^2). \tag{3.1}$$

Then with probability at least $1 - \exp(-cu^2)$ on the random draw of $\{(\mathbf{a}_i, \tau_i)\}_{i=1}^m$, the following holds uniformly for all $\mathbf{x} \in \mathcal{X}$: Let $\mathbf{y} = (y_1, \dots, y_m) \in \mathbb{R}^m$ be given by

$$y_i = \langle \mathbf{a}_i, \mathbf{x} \rangle + \tau_i + v_i, \quad i = 1, \dots, m,$$

such that $(\frac{1}{m} \sum_{i=1}^m v_i^2)^{1/2} \leq \frac{t}{20}$. Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \mathbf{x}\|_2 \leq t$.

Except for the σ -depending summand in (3.1), the above statement does not involve any oversampling factors and it turns into an exact recovery guarantee for $t = 0$ and $\sigma = 0$. Although Corollary 1 bears resemblance with standard results from compressed sensing, it comes with a distinctive feature: The needed number of measurements in (3.1) is determined by the local mean width $w_t(K - \mathcal{X})$. As such, this parameter is quite implicit, so that an informative (upper) bound is required for any specific choice of K and \mathcal{X} . The remainder of this section is devoted to this quest.

Arguably the most popular low-complexity model is *sparsity* in conjunction with an ℓ^1 -relaxation. In our context, this corresponds to the following scenario:

$$K = R\mathbb{B}_1^p \quad \text{and} \quad \mathcal{X} = \{\mathbf{x} \in \mathbb{R}^p : \|\mathbf{x}\|_0 \leq s, \|\mathbf{x}\|_1 = R\}, \tag{3.2}$$

where $s \in [p]$ and $R > 0$; note that the constraint $\|\mathbf{x}\|_1 = R$ is a specific tuning condition for the estimator $(\mathbf{P}_{K,\mathbf{y}})$ that could be further relaxed (see Remark 4(1)). In this special case, one can derive a convenient bound for the local mean width:

$$w_t(K - \mathcal{X}) \leq w_0(K - \mathcal{X}) \lesssim \sqrt{s \cdot \log(2p/s)}. \tag{3.3}$$

Remarkably, this observation dates back to one of the earliest applications of *Gordon’s escape through the mesh theorem* [30] to signal recovery [54]; see Sect. 10 for a proof of (3.3). A combination of (3.3) and Corollary 1 ensures that uniform (exact) reconstruction is feasible with $O(s \cdot \log(2p/s))$ Gaussian measurements. Therefore, our approach is indeed well in line with standard sparse recovery guarantees in compressed sensing, which are typically based on restricted isometry (cf. [21]).

Beyond this classical example, the situation becomes considerably more challenging and forms an important research subject in its own right. A particular obstacle due to uniformity is that $w_t(K - \mathcal{X})$ heavily depends on the size and shape of the signal set \mathcal{X} . On a technical level, we have to deal with a (possibly uncountable) *union of spherical intersections*:

$$w_t(K - \mathcal{X}) = w\left(\bigcup_{\mathbf{x} \in \mathcal{X}} \left[\frac{1}{t}(K - \mathbf{x}) \cap \mathbb{S}^{p-1}\right]\right).$$

Fortunately, this cumbersome expression can be controlled through a much more convenient upper bound, as shown in the following proposition. To the best of our knowledge, this is a new result, which could be of independent interest. Its proof can be found in Sect. 10 and is based on a covering argument.

Proposition 1 *Let $K, \mathcal{X} \subset \mathbb{R}^p$ and $t > 0$. Then,*

$$w_t(K - \mathcal{X}) \lesssim \sup_{\mathbf{x} \in \mathcal{X}} \underbrace{w\left(\frac{1}{2t}(K - \mathbf{x}) \cap \mathbb{B}_2^p\right)}_{=: \tilde{w}_t(K - \mathbf{x})} + t^{-1} \cdot w(\mathcal{X}). \tag{3.4}$$

Moreover, for $\mathbf{x} \in \mathcal{X}$ and K convex, we have that $\tilde{w}_t(K - \mathbf{x}) \lesssim w_0(K - \tilde{\mathbf{x}}) + 1$ for every $\tilde{\mathbf{x}} \in K$ with $\|\mathbf{x} - \tilde{\mathbf{x}}\|_2 \leq t$.

The upper bound in (3.4) implies a significant simplification: The first term is fully *localized*, measuring the complexity of K with respect to each individual point $\mathbf{x} \in \mathcal{X}$,

whereas the second term accounts for the *global* size of \mathcal{X} (independently of the constraint set K). In other words, the effect of K and \mathcal{X} is now decoupled. A second notable fact is that the local mean width $\tilde{w}_t(K - \mathbf{x})$ and its conic counterpart $w_0(K - \mathbf{x})$ are well-known complexity parameters from non-uniform recovery results, e.g., see [2, 12, 23, 51, 63]. Hence, Proposition 1 allows us to transfer any corresponding bound from the literature to the uniform regime. The presence of the global mean width $w(\mathcal{X})$ as an additional expense for uniformity appears natural to us. However, this expression also entails an oversampling factor of t^{-1} , which prevents perfect reconstruction (when applied to Corollary 1). Given the ℓ^1 -special case in (3.3), we suspect that this is an artifact of our proof, and it remains an open question whether such a factor could be removed in general. Our overall conclusion is as follows:

Regarding signal complexity, the transition to uniform recovery requires control over the total size of the set \mathcal{X} , measured by $w(\mathcal{X})$. The constraint set K only appears in terms of local complexity, precisely as in the non-uniform case.

As the above argumentation is fairly abstract, it is insightful to illustrate our approach by a concrete example. In this context, we will also highlight the importance of carefully “designed” signal sets, in the sense that \mathcal{X} (and even its convex hull) is much smaller than K .

Example 1 (Total variation in 1D) We consider the situation of *s-gradient-sparse* signals in one spatial dimension, i.e., for $s \in [p - 1]$ fixed, it is assumed that $\|\nabla \mathbf{x}\|_0 \leq s$, where

$$\nabla := \begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & & 0 \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix} \in \mathbb{R}^{(p-1) \times p}$$

is a discrete gradient operator. Geometrically, this condition simply means that the vector \mathbf{x} is piecewise constant with at most s jump discontinuities. An ℓ^1 -relaxation of gradient sparsity then leads to the common *total variation (TV)* model [55], which in our case corresponds to the constraint set $K = \{\mathbf{x} \in \mathbb{R}^p : \|\nabla \mathbf{x}\|_1 \leq R\}$ for some $R > 0$.

A recent work of the first author [27] has demonstrated that the reconstruction capacity of the TV model does not only depend on the number of jump discontinuities but also strongly on their position. More specifically, we say that an *s-gradient-sparse* signal \mathbf{x} is Δ -separated for some $\Delta \in (0, 1]$ if

$$\min_{j \in [s+1]} \frac{|v_j - v_{j-1}|}{p} \geq \frac{\Delta}{s + 1},$$

where $\text{supp}(\nabla \mathbf{x}) = \{v_1, \dots, v_s\}$ with $0 =: v_0 < v_1 < \dots < v_s < v_{s+1} =: p$. Intuitively, the constant Δ measures how much the jump positions in \mathbf{x} deviate from an equidistant pattern (where one would have $\Delta = 1$). Assuming that $\Delta \in (0, 1]$ is fixed

(independently of n and s), we say that \mathbf{x} is *well separated*; see [27, Sec. 2.1] for more details. This motivates us to consider the following signal set:

$$\mathcal{X} = \{ \mathbf{x} \in \mathbb{R}^p : \|\nabla \mathbf{x}\|_0 \leq s, \mathbf{x} \text{ is } \Delta\text{-separated}, \|\nabla \mathbf{x}\|_1 = R \} \cap \mathbb{B}_2^p. \tag{3.5}$$

Similarly to (3.2), the constraint $\|\nabla \mathbf{x}\|_1 = R$ results from a tuning condition for $(\mathbf{P}_{K,y})$ (see Remark 4(1)).

A deep result from [27, Thm. 2.10] now yields the following estimate for the conic mean width:

$$w_0(K - \mathcal{X}) \lesssim \sqrt{\Delta^{-1} \cdot s \cdot \log^2(p)} \quad \text{for all } \mathbf{x} \in \mathcal{X}. \tag{3.6}$$

For the global mean width, we can make use of a bound derived by Krahmer et al. [38, Sec. 3.2]:

$$w(\mathcal{X}) \lesssim \sqrt{s \cdot \log(ep/s)}. \tag{3.7}$$

Note that (3.7) would also hold without the Δ -separation condition, but it exploits that $\mathcal{X} \subset \mathbb{B}_2^p$. Combining (3.6) and (3.7) with Proposition 1, we finally obtain

$$w_t(K - \mathcal{X}) \lesssim (1 + t^{-1}) \sqrt{s \cdot \log^2(p)}.$$

Therefore, according to Corollary 1, uniform recovery of well-separated s -gradient-sparse signals becomes feasible with $O(s \cdot \log^2(p))$ Gaussian measurements. In particular, we have shown that there is no qualitative gap to the non-uniform guarantee derived in [27].

Let us emphasize that, although the previous bounds look quite simple, their proofs build on fundamental insights into the TV model. For example, relying only on the basic estimate $w_t(K - \mathcal{X}) \lesssim t^{-1} \cdot w(K)$ would not help much, since it is unclear how $w(K)$ scales compared to $w(\mathcal{X})$, cf. [38, Sec. 3.3]. Even more notable is the fact that the above argument would break down if the Δ -separation condition is omitted. Indeed, Cai and Xu [8] have shown that uniform recovery of *all* s -gradient-sparse signals (including pathological examples) via TV minimization is impossible with less than $\Omega(\sqrt{s \cdot p})$ Gaussian measurements. Thus, a striking gap emerges between the uniform and non-uniform regime in this scenario. Remarkably, the technical approach of [8] is based on an adapted nullspace property, for which it is not obvious how to incorporate additional signal structure such as Δ -separation. The presented example on the TV model therefore also underscores the merits of a geometric complexity analysis as proposed in our work. □

While gradient sparsity seems to be only slightly more involved than standard sparsity (with respect to an orthonormal basis), the above consideration has indicated many subtleties. To a certain extent, Example 1 is just the “tip of the iceberg” and similar phenomena apply in the more general context of the analysis- and synthesis- ℓ^1 -model. An in-depth discussion would go beyond the scope of the present paper, and we refer the interested reader to the related works [26, 40]. The actual novelty of Proposition 1

is that these previous findings may also extend to uniform recovery—a step that was not addressed so far.

We close this section with a remark on stable and robust recovery:

Remark 4 (1) *Tuning and stability.* The equality constraint for \mathcal{X} in (3.2) and (3.5) ensures that every $\mathbf{x} \in \mathcal{X}$ lies on the boundary of K . Such a technical step is important for bounds that are based on the *conic* mean width, such as (3.3) and (3.6); if \mathbf{x} would lie in the interior of K , then $\text{cone}(K - \mathbf{x}) = \mathbb{R}^p$ and therefore $w_0(K - \mathbf{x}) \asymp \sqrt{p}$. From an algorithmic perspective, this can be seen as a tuning condition for the Lasso estimator $(P_{K,y})$. The “moreover”-part of Proposition 1 presents a convenient relaxation in this respect: When interested in the local complexity of some point $\mathbf{x} \in \mathcal{X} \subset K$, one may instead consider a nearby point $\tilde{\mathbf{x}}$, whose conic mean width $w_0(K - \tilde{\mathbf{x}})$ is smaller (non-trivial). The price for this simplification is a recovery error in the order of $\|\mathbf{x} - \tilde{\mathbf{x}}\|_2$. This trade-off is closely related to the idea of *stability* (or *compressibility*) in compressed sensing theory [21].

(2) *Robustness and optimality.* The statement of Corollary 1 allows for a comparison of the statistical and adversarial noise model (cf. Remark 1). In the absence of worst-case perturbations (i.e., $v_i = 0$), the condition (3.1) can be translated into an error decay rate of $O(\sigma \cdot m^{-1/2})$. Thus, $(P_{K,y})$ becomes a *consistent* estimator of \mathbf{x} in the ordinary sense of statistics, i.e., $\|\hat{\mathbf{z}} - \mathbf{x}\|_2 \rightarrow 0$ in probability as $m \rightarrow \infty$. Remarkably, the asymptotic error scaling with respect to the noise parameter σ and sample size m cannot be further improved in general—it is minimax optimal for linear observations, e.g., see [52, Sec. 4].

Compared to random noise, the adversarial model is more general, since the v_i may encode any type of perturbation, even deterministic ones. Given the ℓ^2 -constraint in Corollary 1,

$$\left(\frac{1}{m} \sum_{i=1}^m v_i^2\right)^{1/2} \leq \frac{t}{20}, \tag{3.8}$$

it turns out that a recovery error in the order of $O(t)$ is essentially optimal. Indeed, there exist instances of Corollary 1 where the v_i can be selected such that (3.8) holds with high probability (over $\{\mathbf{a}_i\}_{i=1}^m$) and one has that $\|\hat{\mathbf{z}} - \mathbf{x}\|_2 \asymp t$.¹⁰ The difference with the statistical regime becomes clearer when considering random noise from an adversarial perspective, i.e., $v_i = \tau_i \sim N(0, \sigma^2)$. Then, it is not hard to see that with high probability, the constraint (3.8) is only attainable with $t \gtrsim \sigma$, which prevents arbitrarily high precision in terms of t . In particular, Corollary 1 would not certify the consistency of $(P_{K,y})$ anymore, contrary to the argument in the previous paragraph.

Let us emphasize that uniformity does not play a special role in these matters, although the statistical noise model appears somewhat unnatural in this regime

¹⁰ A simple example is as follows: For a unit vector $\mathbf{x} \in \mathbb{S}^{p-1}$ to be reconstructed, set $v_i = c \cdot t \cdot \langle \mathbf{a}_i, \mathbf{x} \rangle$ for a constant $c > 0$ small enough. If m is sufficiently large, then (3.8) holds with high probability. At the same time, we have that $y_i = \langle \mathbf{a}_i, \mathbf{x} \rangle + v_i = \langle \mathbf{a}_i, (1 + ct)\mathbf{x} \rangle$. Hence, Corollary 1 can be also applied such that it certifies exact recovery of the rescaled vector $(1 + ct)\mathbf{x}$ via $(P_{K,y})$. In other words, we have that $\|\hat{\mathbf{z}} - \mathbf{x}\|_2 = ct \cdot \|\mathbf{x}\|_2 \asymp t$.

(see Remark 1). In principle, our conclusions do also carry over to the abstract setting of Theorem 2, but the interpretation of the term “noise” becomes more subtle there. The reason is that the estimator $(P_{K,y})$ basically treats nonlinear distortions as if they would be uncorrelated statistical noise (cf. [51]), even if the output function is completely deterministic. Seen from this angle, better decay rates than $O(\sigma \cdot m^{-1/2})$ are possible, which however strongly depends on the specific model and estimator (see Remark 2(2)).

4 Applications and Examples

This section demonstrates how to derive “out-of-the-box” guarantees from Theorem 2 for specific observation models. Sects. 4.1–4.3 are devoted to applications to quantized compressed sensing. Sect. 4.4 then revisits the case of single-index models, including new applications to modulo measurements and coordinate-wise distortions. Finally, Sect. 4.5 is concerned with a conceptually different example on the problem of variable selection. Note that all proofs are deferred to Sect. 8.

4.1 1-Bit Observations

As already indicated in Sect. 1.1, the most basic version of 1-bit compressed sensing asks for the reconstruction of signals $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^p$ from binary observations $\mathbf{y} \in \{-1, 1\}^m$ of the form

$$y_i = \text{sign}(\langle \mathbf{a}_i, \mathbf{x} \rangle) + v_i, \quad i = 1, \dots, m. \tag{4.1}$$

Here, the noise variables $\mathbf{v} := (v_1, \dots, v_m)$ may take values in $\{-2, 0, 2\}$, modeling possible distortions of the linear measurement process before quantization and/or bit flips during quantization. Importantly, the magnitude of \mathbf{x} gets lost in (4.1) due to the scaling invariance of the sign-function. Thus, the best one can hope for is to reconstruct the direction of \mathbf{x} .

The model of (4.1) is particularly compatible with Gaussian measurement vectors— a related result on sub-Gaussian measurements can be found in the next subsection. Indeed, choosing the target function $T\mathbf{x}$ proportionally to $\mathbf{x}/\|\mathbf{x}\|_2$, one can show that the target mismatch $\rho(\mathbf{x})$ vanishes for every $\mathbf{x} \in \mathcal{X}$. Hence, uniform recovery is possible up to arbitrarily high precision. The following guarantee makes this claim precise and is an application of Theorem 2 to Gaussian 1-bit observations. Its proof in Sect. 8.1 will demonstrate the usefulness of the approximation condition from Assumption 2, see also Fig. 2 for an illustration of the underlying argument.

Corollary 2 *There exist universal constants $c, c_0, C' > 0$ for which the following holds.*

Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent copies of a standard Gaussian random vector $\mathbf{a} \sim N(\mathbf{0}, \mathbf{I}_p)$. Let $\mathcal{X} \subset \mathbb{R}^p$ and define $T\mathbf{x} := \sqrt{\frac{2}{\pi}} \frac{\mathbf{x}}{\|\mathbf{x}\|_2}$ for $\mathbf{x} \in \mathcal{X}$. Moreover, let $K \subset \mathbb{R}^p$ be a convex set such that $T\mathcal{X} \subset K$. For $u \geq 1$ and $t \in (0, 1]$, we assume




Fig. 2 Illustration of the approximation strategy used in the proof of Corollary 2. The basic idea is to approximate the jump discontinuity of $\text{sign}(s)$ (plotted in black) by a linear segment whose slope is inverse proportional to the accuracy t . The resulting function $\psi_t(s)$ (plotted in blue) and the absolute value of the error $\phi_t(s) = |\psi_t(s) - \text{sign}(s)|$ (plotted in dashed red) are then both $\frac{128}{t}$ -Lipschitz. This is already enough to fulfill Assumption 2 with $L_t \lesssim 1 + t^{-1}$ and $\hat{L}_t \lesssim t^{-1}$; note that the factor 128 is just an appropriate constant for the proof (Color figure online)

that

$$m \geq C' \cdot t^{-2} \cdot \left(w_t^2(K - T\mathcal{X}) + t^{-2} \cdot w^2(T\mathcal{X}) + u^2 \right). \tag{4.2}$$

Finally, let $\beta \in [0, 1]$ be such that $\beta\sqrt{\log(e/\beta)} \leq c_0t$. Then with probability at least $1 - \exp(-cu^2)$ on the random draw of $\{\mathbf{a}_i\}_{i=1}^m$, the following holds uniformly for all $\mathbf{x} \in \mathcal{X}$: Let $\mathbf{y} \in \{-1, 1\}^m$ be given by (4.1) such that $\frac{1}{2m}\|\mathbf{v}\|_1 \leq \beta$. Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \sqrt{\frac{2}{\pi}} \frac{\mathbf{x}}{\|\mathbf{x}\|_2}\|_2 \leq t$.

Corollary 2 is in line with some of the early achievements in 1-bit compressed sensing [48, 49]. Remarkably, the condition (4.2) translates into an error decay rate of $O(m^{-1/4})$ in the uniform case, which improves the original rate of $O(m^{-1/12})$ established by Plan and Vershynin in [49, Thm. 1.3]. In fact, we are not aware of any result in the literature that implies the statement of Corollary 2. But we stress that the above oversampling rate is still not optimal (see [34]) and can be further improved with a more specialized argument relying on random hyperplane tessellations (see Corollary 8 in Sect. 5). The noise constraint of Corollary 2 simply means that the fraction of wrong input bits $\frac{1}{2m}\|\mathbf{v}\|_1$ must not exceed β , while the latter may be in the order of t (up to a log-factor). This condition again significantly improves [49, Thm. 1.3] and is a particular consequence of the outlier robustness established in Theorem 2; see also [17] for a similar achievement in the situation of dithered observations.

4.2 1-Bit Observations with Dithering

According to Ai et al. [1], the conclusion of Corollary 2 cannot be extended to sub-Gaussian measurements in general, regardless of the considered reconstruction method. A practicable remedy is the technique of *dithering*, which in its most basic form, corresponds to a random shift of the quantization threshold. Originating from quantized signal processing, e.g., see [13, 31, 32], its benefits recently also emerged in compressed sensing theory [3, 17, 18, 37, 62, 65]. For more background information, we refer to [14] and the references therein.

Extending the original 1-bit model (4.1) by an additional dithering step leads to observations $\mathbf{y} \in \{-1, 1\}^m$ of the following form:

$$y_i = \text{sign}(\langle \mathbf{a}_i, \mathbf{x} \rangle + \tau_i) + v_i, \quad i = 1, \dots, m, \tag{4.3}$$

where $\mathbf{v} := (v_1, \dots, v_m) \in \{-2, 0, 2\}^m$ again models noise. The *dithering variables* τ_i are independent copies of a random variable τ that is uniformly distributed on an interval $[-\lambda, \lambda]$. Note that a major difference between dithering and additive noise is that the parameter $\lambda > 0$ is known and adjustable in practice (while the τ_i could be in principle unknown). The following corollary of Theorem 2 is based on the fact that a careful choice of λ allows us to control the size of the (non-vanishing) target mismatch $\rho(\mathbf{x})$, where $T\mathbf{x} := \lambda^{-1}\mathbf{x}$.

Corollary 3 *There exist universal constants $c, c_0, \tilde{C}, C' > 0$ for which the following holds.*

Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent copies of a centered, isotropic, sub-Gaussian random vector $\mathbf{a} \in \mathbb{R}^p$ with $\|\mathbf{a}\|_{\psi_2} \leq L$. Let τ_1, \dots, τ_m be independent copies of a random variable τ that is uniformly distributed on $[-\lambda, \lambda]$ for a parameter $\lambda > 0$. In addition, suppose that $\{\mathbf{a}_i\}_{i=1}^m$ and $\{\tau_i\}_{i=1}^m$ are independent. Let $\mathcal{X} \subset \mathbb{R}\mathbb{B}_2^p$ for some $R > 0$, and let $K \subset \mathbb{R}^p$ be a convex set such that $\lambda^{-1}\mathcal{X} \subset K$. For $u \geq 1$ and $t \in (0, 1]$, we assume that

$$\begin{aligned} \lambda &\geq \tilde{C} \cdot R \cdot L \cdot \sqrt{\log(e/t)}, \\ m &\geq C' \cdot L^2 \cdot \left((\log L + t^{-2}) \cdot (w_t^2(K - \lambda^{-1}\mathcal{X}) + u^2) + L^2 t^{-4} \lambda^{-2} \cdot w^2(\mathcal{X}) \right). \end{aligned} \tag{4.4}$$

Finally, let $\beta \in [0, 1]$ be such that $\beta \sqrt{\log(e/\beta)} \leq c_0 L^{-2} t$. Then with probability at least $1 - \exp(-cu^2)$ on the random draw of $\{(\mathbf{a}_i, \tau_i)\}_{i=1}^m$, the following holds uniformly for all $\mathbf{x} \in \mathcal{X}$: Let $\mathbf{y} \in \{-1, 1\}^m$ be given by (4.3) such that $\frac{1}{2m} \|\mathbf{y}\|_1 \leq \beta$. Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \lambda^{-1}\mathbf{x}\|_2 \leq t$.

Corollary 3 exhibits many common features of known recovery guarantees based on dithering. In particular, it can be seen as a uniform version of a recent result on the generalized Lasso by Thrampoulidis and Rawat [62, Thm. IV.1]. The most notable improvement over Corollary 2 is that it is now possible to recover the actual signal $\mathbf{x} \in \mathcal{X}$ up to arbitrarily high precision, and not only its direction vector $\mathbf{x}/\|\mathbf{x}\|_2$.

Similarly to Corollary 2, we are not aware of any result in the literature that implies the statement of Corollary 3. Nevertheless, the oversampling rate of $O(m^{-1/4})$ promoted by (4.4) can be further improved with a proof strategy that is specifically tailored to 1-bit observations with dithering (see Corollary 9 in Sect. 5).

4.3 Multi-Bit Observations

While 1-bit measurements are an important extreme case of quantized compressed sensing, a considerable part of the literature deals with multi-bit observation models; once again, see [6, 14] for a good introduction to this subject. In this subsection, we illustrate our approach in the prototypical situation of *uniform quantization*: For a fixed $\delta > 0$, the goal is to reconstruct a signal $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^p$ from observations $\mathbf{y} \in \delta\mathbb{Z}^m$ of the form

$$y_i = q_\delta(\langle \mathbf{a}_i, \mathbf{x} \rangle + \tau_i) + v_i, \quad i = 1, \dots, m, \tag{4.5}$$

where $q_\delta(v) := (2\lceil \frac{v}{2\delta} \rceil - 1)\delta$ is a uniform quantizer on the grid $\delta(2\mathbb{Z} - 1)^m$ with resolution $\delta > 0$. The dithering variables τ_i are independent copies of a random variable τ that is uniformly distributed on $[-\delta, \delta]$. As before, $\mathbf{v} := (v_1, \dots, v_m)$ can describe any type of adversarial noise, but now takes values in $\delta\mathbb{Z}^m$. Theorem 2 yields the following uniform recovery guarantee for multi-bit observations with sub-Gaussian measurements and dithering.

Corollary 4 *There exist universal constants $c, c_0, C' > 0$ for which the following holds.*

Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent copies of a centered, isotropic, sub-Gaussian random vector $\mathbf{a} \in \mathbb{R}^p$ with $\|\mathbf{a}\|_{\psi_2} \leq L$. Let τ_1, \dots, τ_m be independent copies of a random variable τ that is uniformly distributed on $[-\delta, \delta]$ for a fixed parameter $\delta > 0$. In addition, suppose that $\{\mathbf{a}_i\}_{i=1}^m$ and $\{\tau_i\}_{i=1}^m$ are independent. Let $\mathcal{X} \subset \mathbb{R}^p$ be a bounded subset, and let $K \subset \mathbb{R}^p$ be a convex set such that $\mathcal{X} \subset K$. For $u \geq 1$ and $t > 0$, we assume that

$$m \geq C' \cdot L^2 \cdot \left((\log L + t^{-2}\delta^2) \cdot (w_t^2(K - \mathcal{X}) + u^2) + L^2 t^{-4} \delta^2 \cdot w^2(\mathcal{X}) \right). \tag{4.6}$$

Then with probability at least $1 - \exp(-cu^2)$ on the random draw of $\{(\mathbf{a}_i, \tau_i)\}_{i=1}^m$, the following holds uniformly for all $\mathbf{x} \in \mathcal{X}$: Let $\mathbf{y} \in \delta\mathbb{Z}^m$ be given by (4.5) such that at least one of the following conditions is fulfilled:

$$\begin{aligned} (a) \quad & \frac{1}{\sqrt{m}} \|\mathbf{v}\|_2 \leq \frac{t}{40} \quad \text{or} \\ (b) \quad & \frac{1}{\sqrt{m}} \|\mathbf{v}\|_2 \leq \frac{c_0 \sqrt{\delta t}}{L^2 \sqrt{\max\{1, \log(\delta e/t)\}}} \quad \text{and} \quad \frac{1}{m} \|\mathbf{v}\|_0 \leq \frac{t}{8}. \end{aligned} \tag{4.7}$$

Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \mathbf{x}\|_2 \leq 2t$.

Corollary 4 can be considered as a uniform version of a recent result by Thrampoulidis and Rawat [62, Thm. III.1]. For a fixed quantizer resolution $\delta > 0$, the condition (4.6)

translates into an error decay rate of $O(\sqrt{\delta} \cdot m^{-1/4})$. On the other hand, if $\delta \ll t$, then (4.6) is satisfied as soon as $m \geq C \cdot L^2 \log L \cdot (w_t^2(K - \mathcal{X}) + u^2)$. In other words, if the quantizer resolution is much higher than the desired reconstruction accuracy, the performance of $(P_{K,y})$ is the same as if the input vector would consist of linear measurements (cf. Corollary 1). This behavior is perfectly consistent with the fact that the uniformly quantized observations (4.5) become linear as $\delta \rightarrow 0$. It is also worth pointing out that the constraint (b) in (4.7) implies a certain outlier robustness: If $t < \delta$ and only a fraction of t/δ bits are corrupted, then the normalized ℓ^2 -noise error may scale in the order of $\sqrt{\delta t}$ instead of t .

To the best of our knowledge, Corollary 4 is a new result. But we stress that there exist uniform recovery guarantees for other programs than $(P_{K,y})$ in the literature, which imply better oversampling rates, e.g., see [36, Thm. 3] or [65] in the case of structured signal sets. Nevertheless, we expect that Corollary 4 could be easily improved in that regard, using the strategy of Sect. 5 in conjunction with known uniform embedding results.

4.4 Single-Index Models and Beyond

The situation of single-index models was already considered in Theorem 1 in Sect. 1.2, as an appetizer for our more general approach. As pointed out there, this guarantee can be seen as an upgrade of an earlier non-uniform result by Plan and Vershynin [51, Thm. 1.9], thereby demonstrating that uniform recovery is possible beyond quantized measurement schemes. In view of Assumption 2, the Lipschitz continuity of f in Theorem 1 is certainly not necessary. However, it verifies that a more “regular” observation variable can lead to a near-optimal error decay rate (see Remark 2(2)). Apart from that, we emphasize that (Lipschitz) continuous nonlinearities are not only of academic interest but also appear in practical applications, for instance, as power amplifiers in sensor networks [24].

The remainder of this subsection is devoted to two new applications of our general framework, which extend Theorem 1 into different directions.

Modulo measurements. The first scenario is inspired by a recent work of Bhandari et al. [4] on “unlimited sampling.” The practical motivation there is to prevent a saturation of analog-to-digital converters by applying a modulo operator in the measurement process. Tailored to the setup of the present article, we consider *modulo measurements* of the form

$$y_i = m_\lambda(\langle a_i, x \rangle) + v_i, \quad i = 1, \dots, m, \tag{4.8}$$

where the modulo function is given by $m_\lambda(v) := v - \lfloor \frac{v+\lambda}{2\lambda} \rfloor \cdot 2\lambda$ for a fixed parameter $\lambda > 0$, and $v := (v_1, \dots, v_m) \in \mathbb{R}^m$ models noise. Clearly, the nonlinearity m_λ neither corresponds to a quantization (cf. Sects. 4.1–4.3) nor is it Lipschitz continuous (cf. Theorem 1). The following corollary of Theorem 2 provides a uniform recovery guarantee for the observation model (4.8):

Corollary 5 *There exist universal constants $c, C, C' > 0$ for which the following holds.*

Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent copies of a standard Gaussian random vector $\mathbf{a} \sim N(\mathbf{0}, \mathbf{I}_p)$. For $\lambda > 0$, we set $\mu_\lambda := \mathbb{E}[\mathfrak{m}_\lambda(g)g]$ with $g \sim N(0, 1)$. Let $\mathcal{X} \subset \mathbb{S}^{p-1}$, and let $K \subset \mathbb{R}^p$ be a convex set such that $\mu_\lambda \mathcal{X} \subset K$. For $u \geq 1$ and $t \in (0, \lambda]$, we assume that $\lambda \geq C$ and

$$m \geq C' \cdot \left((1 + t^{-2}) \cdot (w_t^2(K - \mu_\lambda \mathcal{X}) + u^2) + \lambda^2 t^{-4} \cdot w^2(\mathcal{X}) \right).$$

Then, $\mu_\lambda \in [\frac{1}{2}, 1]$ and with probability at least $1 - \exp(-cu^2)$ on the random draw of $\{\mathbf{a}_i\}_{i=1}^m$, the following holds uniformly for all $\mathbf{x} \in \mathcal{X}$: Let $\mathbf{y} \in \mathbb{R}^m$ be given by (4.8) such that $(\frac{1}{m} \sum_{i=1}^m v_i^2)^{1/2} \leq \frac{1}{20}$. Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \mu_\lambda \mathbf{x}\|_2 \leq t$.

Similar to our above applications to quantized compressed sensing, the (infinitely many) discontinuities of the modulo function \mathfrak{m}_λ imply an error decay rate of $O(m^{-1/4})$. To the best of our knowledge, Corollary 5 is a new result in its own right. In particular, we are not aware of alternative proof techniques that would allow us to derive a comparable statement. Finally, it is worth noting that the condition $\lambda \geq C$ in Corollary 5 could be further relaxed to $\lambda > 0$. This would come at the cost of a lower bound for μ_λ , since we have that $\mu_\lambda \rightarrow 0$ as $\lambda \rightarrow 0$ (the nonlinearity \mathfrak{m}_λ uniformly converges to 0 as $\lambda \rightarrow 0$).

Coordinate-wise distortions. The common single-index model (see Theorem 1) assumes that a nonlinear output function perturbs the linear observations $\langle \mathbf{a}_i, \mathbf{x} \rangle$. Instead, one can also imagine distortions that affect the computation of the dot product $\langle \mathbf{a}_i, \mathbf{x} \rangle$ directly. More specifically, we are interested in observations of the form

$$y_i = f(\mathbf{a}_i \circ \mathbf{x}) + v_i, \quad i = 1, \dots, m, \tag{4.9}$$

where $\mathbf{a} \circ \mathbf{b} \in \mathbb{R}^p$ denotes the coordinate-wise product of vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^p$ and $f : \mathbb{R}^p \rightarrow \mathbb{R}$ can be represented as

$$f(\mathbf{z}) := \sum_{j=1}^p f_j(z_j), \quad \mathbf{z} \in \mathbb{R}^p.$$

Here, we assume that the functions $f_j : \mathbb{R} \rightarrow \mathbb{R}$ are odd and γ -Lipschitz, while satisfying the growth conditions

$$\alpha(v - v') \leq f_j(v) - f_j(v'), \quad \text{for all } v, v' \in \mathbb{R} \text{ with } v' \leq v, \tag{4.10}$$

and

$$\beta_1 v \leq f_j(v) \leq \beta_2 v, \quad \text{for all } v \geq 0, \tag{4.11}$$

with parameters $\alpha > 0$ and $\beta_2 \geq \beta_1 > 0$. One can think of (4.9) as computing a “nonlinear dot product” of \mathbf{a}_i and \mathbf{x} . While models of this type are not covered by the

original approach of [51], an application of Theorem 2 allows us to deal with them under appropriate assumptions:

Corollary 6 *There exist universal constants $c, C > 0$ for which the following holds.*

Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent copies of a centered, isotropic random vector $\mathbf{a} = (a_1, \dots, a_p)$ which has independent, symmetric, and sub-Gaussian coordinates such that $\max_{j \in [p]} \|a_j\|_{\psi_2} \leq L$ for some $L > 0$. Let $\mathcal{X} \subset \mathbb{R}\mathbb{B}_2^p$ and define $T : \mathcal{X} \rightarrow \mathbb{R}^p$ by $T\mathbf{x} := \mathbb{E}[f(\mathbf{a} \circ \mathbf{x})\mathbf{a}]$. Let $K \subset \mathbb{R}^p$ be a convex set such that $T\mathcal{X} \subset K$. For $u \geq 1$ and $t \geq 0$, we assume that

$$m \geq C \cdot L^4 \cdot \left((1 + t^{-2}R^2(\beta_2^2 + \gamma^2)) \cdot (w_t^2(K - T\mathcal{X}) + u^2) + t^{-2}(\frac{\gamma}{\alpha})^2 \cdot w^2(T\mathcal{X}) \right).$$

Then with probability at least $1 - \exp(-cu^2)$ on the random draw of $\{\mathbf{a}_i\}_{i=1}^m$, the following holds uniformly for all $\mathbf{x} \in \mathcal{X}$: Let $\mathbf{y} \in \mathbb{R}^m$ be given by (4.9) such that $(\frac{1}{m} \sum_{i=1}^m v_i^2)^{1/2} \leq \frac{t}{20}$. Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - T\mathbf{x}\|_2 \leq t$. Furthermore, for every $j \in [p]$, we have that

$$\begin{aligned} \beta_1 x_j &\leq (T\mathbf{x})_j \leq \beta_2 x_j, & \text{for } x_j \geq 0, \\ \beta_2 x_j &\leq (T\mathbf{x})_j \leq \beta_1 x_j, & \text{for } x_j < 0. \end{aligned} \tag{4.12}$$

The above result shows that estimation via $(P_{K,\mathbf{y}})$ is even possible in situations where instead of linear measurements $\langle \mathbf{a}_i, \mathbf{x} \rangle = \sum_{j=1}^p (\mathbf{a}_i)_j x_j$ one observes distorted dot products $f(\mathbf{a}_i \circ \mathbf{x}) = \sum_{j=1}^p f_j((\mathbf{a}_i)_j x_j)$. Note that the generalized Lasso may not precisely recover a scalar multiple of \mathbf{x} here, but rather a vector $T\mathbf{x}$ that lies in a hyperrectangle defined by \mathbf{x} and the parameters β_1 and β_2 . The closer the functions f_j are to the identity, the closer β_1 and β_2 are to 1, which implies that $T\mathbf{x} \approx \mathbf{x}$ according to (4.12).

4.5 Variable Selection

This subsection is devoted to an instance of Assumption 1 that is substantially different from the previous ones: For a fixed integer $s \leq p$ and an output function $f : \mathbb{R}^p \rightarrow \mathbb{R}$, we ask for estimation of an index set $\mathcal{S} \subset [p]$ with $|\mathcal{S}| \leq s$ from observations $\mathbf{y} \in \mathbb{R}^m$ of the form

$$y_i = f(\mathbf{a}_{i,\mathcal{S}}) + v_i, \quad i = 1, \dots, m. \tag{4.13}$$

Here, $\mathbf{a}_{i,\mathcal{S}} \in \mathbb{R}^p$ is the coordinate projection of \mathbf{a}_i onto \mathcal{S} and $v_i \in \mathbb{R}$ models additive noise. Most notably, the signals of interest are not parameters vectors in \mathbb{R}^p anymore but correspond to (small) index sets, specifying those coefficients of a measurement vector that contribute to the observation. From a statistical perspective, this can be seen as a *variable selection model*, where $\mathcal{S} \subset [p]$ determines the set of active variables among all feature variables in \mathbf{a}_i . In the context of uniform recovery, this leads to the following problem: Given a collection of sample data $\{\mathbf{a}_i\}_{i=1}^m$, can we retrieve any possible index set $\mathcal{S} \subset [p]$ with $|\mathcal{S}| \leq s$ from (nonlinear) observations of the form

(4.13)? The following corollary of Theorem 2 provides a result in that direction under natural conditions on the output function f .

Corollary 7 *There exist universal constants $c, C > 0$ for which the following holds.*

Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent copies of a centered, isotropic random vector $\mathbf{a} = (a_1, \dots, a_p)$ which has independent sub-Gaussian coordinates such that $\max_{j \in [p]} \|a_j\|_{\psi_2} \leq L$ for some $L > 0$. For $s \leq p$ and an output function $f : \mathbb{R}^p \rightarrow \mathbb{R}$, let $\mathcal{X} \subset \{\mathcal{S} \subset [p] : |\mathcal{S}| \leq s\}$ and define $T : \mathcal{X} \rightarrow \mathbb{R}^p$ by $T\mathcal{S} := \mathbb{E}[f(\mathbf{a}_{\mathcal{S}})\mathbf{a}]$, where $\mathbf{a}_{\mathcal{S}} \in \mathbb{R}^p$ denotes the coordinate projection of \mathbf{a} onto \mathcal{S} . Moreover, we assume that there exist parameters $\alpha, \beta, \gamma, \kappa > 0$ such that

$$\frac{\alpha}{\sqrt{s}} \leq |(T\mathcal{S})_j| \leq \frac{\beta}{\sqrt{s}} \quad \text{for all } \mathcal{S} \in \mathcal{X} \text{ and } j \in \mathcal{S} \tag{4.14}$$

and

$$\|f(\mathbf{a}_{\mathcal{S}}) - f(\mathbf{a}_{\mathcal{S}'})\|_{\psi_2} \leq \gamma \sqrt{\frac{|\mathcal{S} \Delta \mathcal{S}'|}{s}} \quad \text{and} \quad \|f(\mathbf{a}_{\mathcal{S}})\|_{\psi_2} \leq \kappa \quad \text{for all } \mathcal{S}, \mathcal{S}' \in \mathcal{X}, \tag{4.15}$$

where $\mathcal{S} \Delta \mathcal{S}' \subset [p]$ denotes the symmetric difference between \mathcal{S} and \mathcal{S}' . Let $K \subset \mathbb{R}^p$ be a convex set such that $T\mathcal{X} \subset K$. For $u \geq 1$ and $t \geq 0$, we assume that

$$m \geq C \cdot L^2 \cdot \left((\log L + t^{-2}(L^2\beta^2 + \kappa^2)) \cdot (w_t^2(K - T\mathcal{X}) + u^2) + t^{-2}(L^2 + \gamma^2\alpha^{-2}) \cdot w^2(T\mathcal{X}) \right). \tag{4.16}$$

Then with probability at least $1 - \exp(-cu^2)$ on the random draw of $\{\mathbf{a}_i\}_{i=1}^m$, the following holds uniformly for all $\mathcal{S} \in \mathcal{X}$: Let $\mathbf{y} \in \mathbb{R}^m$ be given by (4.13) such that $(\frac{1}{m} \sum_{i=1}^m v_i^2)^{1/2} \leq \frac{t}{20}$. Then, $\text{supp}(T\mathcal{S}) = \mathcal{S}$ and every minimizer $\hat{\mathbf{z}}$ of $(P_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - T\mathcal{S}\|_2 \leq t$.

The assumptions (4.14) and (4.15) can be seen as natural balancing properties of the underlying observation model: (4.14) requires that the coefficients of each target vector $T\mathcal{S} \in K$ are uniformly bounded below and above on \mathcal{S} (and in particular $\alpha \leq \|T\mathcal{S}\|_2 \leq \beta$). The increment condition (4.15) ensures that the distance between two observation variables can be controlled in terms of the symmetric difference of their associated index sets in \mathcal{X} .

With this in mind, Corollary 7 suggests the following simple procedure for variable selection: First perform a hard-thresholding step on $\hat{\mathbf{z}}$ to extract its largest entries in magnitude; then use the corresponding indices to estimate the set of active variables $\mathcal{S} = \text{supp}(T\mathcal{S})$. Note that this does not require explicit knowledge of the output function f . However, the (guaranteed) success of such a strategy strongly depends on the size of α and the accuracy t . In the worst case, t would have to be in the order of α/\sqrt{s} for perfect recovery of \mathcal{S} , which would lead to an undesirable factor of s in (4.16). It is certainly possible to show refined versions of Corollary 7. We suspect that

sharper error bounds could be obtained by considering a different error measure than the ℓ^2 -norm (see Remark 3(3)). Nevertheless, a detailed elaboration would go beyond the scope of this article, and we confine ourselves with the above proof of concept.

5 Uniform Recovery Without Increment Conditions

We have seen in Sects. 4.1–4.3 that the increment conditions of Assumption 2 can lead to inferior error decay rates for quantizing output functions, due to their points of discontinuity. In this section, we present a workaround that can do without any (sub-Gaussian) increment conditions and thereby enables significantly better rates. The basic idea is to cover the signal set \mathcal{X} by an ε -net \mathcal{X}_ε for some $\varepsilon > 0$ and to apply the *non-uniform* version of Theorem 2 to each $\mathbf{x} \in \mathcal{X}_\varepsilon$ separately. Taking the union bound then yields uniform recovery on \mathcal{X}_ε . The final, and most crucial, ingredient that allows us to pass over to the entire signal set \mathcal{X} is the following *local stability condition* on the observation model. It is based on the (outlier) noise bounds in (2.4) and essentially requires that close points in $T\mathcal{X}$ imply close observation vectors.

Assumption 3 Let Assumption 1 be satisfied, and let $t > 0$ and $\varepsilon > 0$. For $m_0 \in \{0, 1, \dots, \lfloor \frac{m}{2} \rfloor\}$, $\Delta > 0$, and $\eta \in [0, 1]$, we assume that the following holds with probability at least $1 - \eta$:

$$\begin{aligned} \sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \|T\mathbf{x} - T\mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{\sqrt{m}} \|\tilde{\mathbf{y}}(\mathbf{x}) - \tilde{\mathbf{y}}(\mathbf{x}')\|_{[2m_0]} &\leq \frac{1}{2} \Delta t \quad \text{and} \\ \sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \|T\mathbf{x} - T\mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{\sqrt{m}} \sigma_{m_0}(\tilde{\mathbf{y}}(\mathbf{x}) - \tilde{\mathbf{y}}(\mathbf{x}'))_2 &\leq \frac{t}{40}. \end{aligned} \tag{5.1}$$

The above strategy leads to the following general uniform recovery guarantee; see Sect. 9 for a detailed proof.

Theorem 3 *There exist universal constants $c, C, C_0 > 0$ for which the following holds.*

Let Assumptions 1 and 3 be satisfied, let $r := \sup_{\mathbf{x} \in \mathcal{X}} \|\langle \mathbf{a}, T\mathbf{x} \rangle - \tilde{\mathbf{y}}(\mathbf{x})\|_{\psi_2}$, and assume $\rho(\mathbf{x}) \leq \frac{t}{32}$ for every $\mathbf{x} \in \mathcal{X}$. For $u \geq 1$ and $u_0 \geq \sqrt{2m_0 \log(em/2m_0)}$, we assume that

$$\begin{aligned} m \geq C \cdot L^2 \cdot \left((\log L + t^{-2}r^2) \cdot \left(\sup_{\mathbf{x} \in \mathcal{X}} w_r^2(K - T\mathbf{x}) + u^2 \right) \right. \\ \left. + L^2 \Delta^2 \cdot \left(\sup_{\mathbf{x} \in \mathcal{X}} w_r^2(K - T\mathbf{x}) + u_0^2 \right) \right) \end{aligned} \tag{5.2}$$

and

$$\min\{u^2, u_0^2\} \geq C_0 \cdot \log \mathcal{N}(T\mathcal{X}, \varepsilon). \tag{5.3}$$

Then with probability at least $1 - \exp(-cu^2) - \exp(-cu_0^2) - \eta$ on the random draw of $\{\langle \mathbf{a}_i, F_i \rangle\}_{i=1}^m$, the following holds uniformly for every $\mathbf{x} \in \mathcal{X}$: Let $\mathbf{y} \in \mathbb{R}^m$ be any

input vector such that

$$\frac{1}{\sqrt{m}} \| \mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}) \|_{[2m_0]} \leq \frac{1}{2} \Delta t \quad \text{and} \quad \frac{1}{\sqrt{m}} \sigma_{m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}))_2 \leq \frac{t}{40}.$$

Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K, \mathbf{y}})$ satisfies $\| \hat{\mathbf{z}} - T\mathbf{x} \|_2 \leq t + \varepsilon$.

The statement of Theorem 3 strongly resembles the one of Theorem 2 with the important difference that the former does not rely on Assumption 2. In particular, the condition (5.2) does not depend on the increment parameters L_t and \hat{L}_t anymore, making it less restrictive than (2.3). It is worth pointing out that the global complexity of \mathcal{X} is now measured in terms of the covering number $\mathcal{N}(T\mathcal{X}, \varepsilon)$ in (5.3) instead of the mean width $w(T\mathcal{X})$. Furthermore, (5.2) establishes a refined local complexity measure, due to $\sup_{\mathbf{x} \in \mathcal{X}} w_t^2(K - T\mathbf{x}) \leq w_t^2(K - T\mathcal{X})$; see also Proposition 1. Nevertheless, the gain of Theorem 3 is obviously linked to the verification of Assumption 3, which is usually a highly non-trivial task.

We now present two applications of Theorem 3 in the prototypical case of 1-bit observations. In this specific situation, it turns out that Assumption 3 is compatible with uniform bounds for binary embeddings, which allows us to make use of related results from the literature. Our first application is based on the following embedding guarantee by Oymak and Recht [45] for noiseless, Gaussian 1-bit measurements (cf. Sect. 4.1). Hereafter, we agree on the shortcut notation $[H]_\varepsilon := (\frac{1}{\varepsilon}H - \frac{1}{\varepsilon}H) \cap \mathbb{B}_2^p$ for any subset $H \subset \mathbb{R}^p$ and $\varepsilon > 0$; also note that the parameter $w([H]_\varepsilon)$ is virtually the same as the local mean width $w_\varepsilon(H - H)$ except that the former intersects with \mathbb{B}_2^p instead of \mathbb{S}^{p-1} .

Theorem 4 ([45, Thm. 3.2]) *There exist universal constants $c, \bar{c}, C > 0$ for which the following holds.*

Let $H \subset \mathbb{S}^{p-1}$, and let $\mathbf{A} \in \mathbb{R}^{m \times p}$ be a random matrix with independent standard Gaussian row vectors $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$. For $\beta \in (0, 1)$ and $\varepsilon \leq \bar{c}\beta/\sqrt{\log(e/\beta)}$, we assume that

$$m \geq C \cdot \left(\varepsilon^2 \beta^{-3} \cdot w^2([H]_\varepsilon) + \beta^{-1} \cdot \log \mathcal{N}(H, \varepsilon) \right).$$

Then, the following holds with probability at least $1 - \exp(-cm\beta)$:

$$\sup_{\substack{\mathbf{x}, \mathbf{x}' \in H \\ \|\mathbf{x} - \mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{2m} \| \text{sign}(\mathbf{A}\mathbf{x}) - \text{sign}(\mathbf{A}\mathbf{x}') \|_1 \leq \beta.$$

Combining Theorem 4 with Theorem 3 yields an improved version of Corollary 2 (see Sect. 9 for a proof):

Corollary 8 *There exist universal constants $c, c', c_0, C', C_0 > 0$ for which the following holds.*

Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent copies of a standard Gaussian random vector $\mathbf{a} \sim N(\mathbf{0}, \mathbf{I}_p)$. Let $\mathcal{X} \subset \mathbb{R}^p$ and define $T\mathbf{x} := \sqrt{\frac{2}{\pi}} \frac{\mathbf{x}}{\|\mathbf{x}\|_2}$ for $\mathbf{x} \in \mathcal{X}$. Moreover, let

$K \subset \mathbb{R}^p$ be a convex set such that $T\mathcal{X} \subset K$. Fix $t \in (0, 1)$ and $\varepsilon \leq c't / \log(e/t)$. For $u^2 \geq \max\{1, C_0 \cdot \log \mathcal{N}(T\mathcal{X}, \varepsilon)\}$, we assume that

$$m \geq C' \cdot \left(t^{-2} \cdot \left(\sup_{\mathbf{x} \in \mathcal{X}} w_t^2(K - T\mathbf{x}) + u^2 \right) + \varepsilon^2 t^{-3} \log^{3/2}(e/t) \cdot w^2([T\mathcal{X}]_\varepsilon) \right). \tag{5.4}$$

Finally, let $\beta \in [0, 1]$ be such that $\beta \sqrt{\log(e/\beta)} \leq c_0 t$. Then with probability at least $1 - \exp(-cu^2)$ on the random draw of $\{\mathbf{a}_i\}_{i=1}^m$, the following holds uniformly for all $\mathbf{x} \in \mathcal{X}$: Let $\mathbf{y} \in \{-1, 1\}^m$ be given by (4.1) such that $\frac{1}{2m} \|\mathbf{y}\|_1 \leq \frac{\beta}{2}$. Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \sqrt{\frac{2}{\pi}} \frac{\mathbf{x}}{\|\mathbf{x}\|_2}\|_2 \leq 2t$.

To better understand the oversampling behavior of this result, let us consider the situation where $\varepsilon \asymp t / \log(e/t)$. Due to Sudakov minoration, we have that $\log \mathcal{N}(T\mathcal{X}, \varepsilon) \lesssim \varepsilon^{-2} \cdot w^2(T\mathcal{X})$ (e.g., see [64, Thm. 7.4.1]). Hence, in the *worst case*, Corollary 8 would imply an error decay rate of $O(m^{-1/4})$ up to log factors. However, for many low-dimensional signal sets, such as subspaces or sparse vectors, it is possible to establish a much stronger bound of the form $\log \mathcal{N}(T\mathcal{X}, \varepsilon) \lesssim \log(\varepsilon^{-1}) \cdot w^2(T\mathcal{X})$, e.g., see [45, Sec. 2].¹¹ In this case, Corollary 8 would yield an error decay rate of $O(m^{-1/2})$ up to log factors. This is superior to the achievement of Corollary 2 and matches the best possible rate that can be expected for $(P_{K,\mathbf{y}})$ in general (see Remark 2(2)).

Our second application of Theorem 3 concerns the setup of sub-Gaussian 1-bit observations with dithering, as considered in Sect. 4.2. In this case, the stability condition of Assumption 3 can be related to a recent embedding result of Dirksen and Mendelson [17], which is based on hyperplane tessellations for sub-Gaussian vectors.

Theorem 5 ([17, Thm. 2.9]) *For every $L > 0$, there exist constants $c, \bar{c}, C, \tilde{C} > 0$ only depending on L for which the following holds.*

Let $\mathcal{X} \subset \mathbb{R}\mathbb{B}_2^p$ for some $R > 0$, and let $\mathbf{A} \in \mathbb{R}^{m \times p}$ be a random matrix with independent, isotropic, sub-Gaussian row vectors $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$. Moreover, assume that $\max_{i \in [m]} \|\mathbf{a}_i\|_{\psi_2} \leq L$. Let $\boldsymbol{\tau} := (\tau_1, \dots, \tau_m) \in \mathbb{R}^m$ be a random vector with independent entries that are uniformly distributed on $[-\lambda, \lambda]$ for a parameter $\lambda \geq \tilde{C} \cdot R$. In addition, suppose that \mathbf{A} and $\boldsymbol{\tau}$ are independent. For $\beta \in (0, 1]$ and $\varepsilon \leq \bar{c}\beta / \sqrt{\log(e/\beta)}$, we assume that

$$m \geq C \cdot \left(\varepsilon^2 \beta^{-3} \cdot w^2([(2\lambda)^{-1}\mathcal{X}]_\varepsilon) + \beta^{-1} \cdot \log \mathcal{N}(\lambda^{-1}\mathcal{X}, \varepsilon) \right).$$

Then, the following holds with probability at least $1 - \exp(-cm\beta)$:

$$\sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \lambda^{-1} \|\mathbf{x} - \mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{2m} \|\text{sign}(\mathbf{A}\mathbf{x} + \boldsymbol{\tau}) - \text{sign}(\mathbf{A}\mathbf{x}' + \boldsymbol{\tau})\|_1 \leq \beta.$$

¹¹ Here, it is important to note that, in contrast to K , the (transformed) signal set $T\mathcal{X}$ may be highly non-convex.

Combining Theorem 5 with Theorem 3 now yields an improved version of Corollary 3 (see Sect. 9 for a proof):

Corollary 9 *For every $L > 0$, there exist universal constants $c_0, C_0 > 0$ and constants $c, c', \tilde{C}, C' > 0$ only depending on L for which the following holds.*

Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent copies of a centered, isotropic, sub-Gaussian random vector $\mathbf{a} \in \mathbb{R}^p$ with $\|\mathbf{a}\|_{\psi_2} \leq L$. Let τ_1, \dots, τ_m be independent copies of a random variable τ that is uniformly distributed on $[-\lambda, \lambda]$ for a parameter $\lambda > 0$. In addition, suppose that $\{\mathbf{a}_i\}_{i=1}^m$ and $\{\tau_i\}_{i=1}^m$ are independent. Let $\mathcal{X} \subset \mathbb{R}\mathbb{B}_2^p$ for some $R > 0$, and let $K \subset \mathbb{R}^p$ be a convex set such that $\lambda^{-1}\mathcal{X} \subset K$. Fix $t \in (0, 1)$ and $\varepsilon \leq c't/\log(e/t)$. For $u^2 \geq \max\{1, C_0 \cdot \log \mathcal{N}(\lambda^{-1}\mathcal{X}, \varepsilon)\}$, we assume that

$$\lambda \geq \tilde{C} \cdot R \cdot \sqrt{\log(e/t)},$$

$$m \geq C' \cdot \left(t^{-2} \cdot \left(\sup_{\mathbf{x} \in \mathcal{X}} w_t^2(K - \lambda^{-1}\mathbf{x}) + u^2 \right) + \varepsilon^2 t^{-3} \log^{3/2}(e/t) \cdot w^2([\!(2\lambda)^{-1}\mathcal{X}]_\varepsilon) \right). \tag{5.5}$$

Finally, let $\beta \in [0, 1]$ be such that $\beta\sqrt{\log(e/\beta)} \leq c_0L^{-2}t$. Then with probability at least $1 - \exp(-cu^2)$ on the random draw of $\{(\mathbf{a}_i, \tau_i)\}_{i=1}^m$, the following holds uniformly for all $\mathbf{x} \in \mathcal{X}$: Let $\mathbf{y} \in \{-1, 1\}^m$ be given by (4.3) such that $\frac{1}{2m}\|\mathbf{y}\|_1 \leq \frac{\beta}{2}$. Then, every minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \lambda^{-1}\mathbf{x}\|_2 \leq 2t$.

Analogously to noiseless 1-bit observations, Corollary 9 implies a decay rate of $O(m^{-1/4})$ up to log factors in the worst case, which may be improved to $O(m^{-1/2})$ for structured signal sets. This is in line with a recent finding of Dirksen and Mendelson [17, Thm. 1.7], who have analyzed a different estimator.

We close this section by pointing out that the error bounds achieved by Corollary 8 and 9 do still not match the information-theoretic optimal rate of $O(m^{-1})$, e.g., see [18, 34]. We suspect that this gap is not an artifact of our proof, but rather due to a fundamental performance limit of the generalized Lasso $(P_{K,\mathbf{y}})$. A potential remedy is to consider a more specialized reconstruction method, e.g., see [36].

6 Conclusion

This section highlights several key aspects of this work and discusses them in the light of our initial objectives as well as remaining challenges.

- (1) *Observation models.* Probably the greatest benefit of our approach to Problem 1 is its flexibility. The setup of Assumption 1 allows us to implement and analyze virtually every nonlinear observation model that is conceivable for the program $(P_{K,\mathbf{y}})$. The examples that we have seen in Sect. 4 are only a selection of possible applications. Importantly, each of these results is accompanied by a conceptually simple proof that does not require any deeper insight into the underlying model mechanisms; see also the proof template at the beginning of Sect. 8. Therefore, our methodology may form a general path toward competitive benchmark guarantees. However, the resulting oversampling rates do not always match the best possible

rates from the literature, due to the increment condition of Assumption 2. In Sect. 5, we have presented a workaround for this issue, based on an additional local stability assumption. This strategy can indeed lead to near-optimal error bounds, but it relies on the availability of a strong embedding result in each considered case.

- (2) *Non-Gaussian measurements.* Our main result reveals the impact of non-Gaussian measurements, in particular, when consistent recovery via $(P_{K,y})$ can be expected and when not. While relaxing the conditions on isotropy and sub-Gaussian tails in Assumption 1 is certainly practicable (see Remark 3(1)), our analysis is inherently limited to *independent* measurement vectors. This excludes more structured measurement ensembles, such as partial random circulant matrices or lossless expanders [21, Chap. 12 & 13]. On the other hand, such schemes typically satisfy a variant of the restricted isometry property, which is the basis for some recent advances in quantized compressed sensing, e.g., see [16, 20, 33, 65]. These findings provide evidence that at least parts of our results remain valid for a much larger family of measurement vectors. Thus, any guarantee with structured measurements that meets the generality of Theorem 2 would be significant.
- (3) *Complexity parameters.* The research on high-dimensional signal estimation has shown that the (local) Gaussian mean width is a natural, yet accurate complexity measure for convex programs such as $(P_{K,y})$. A distinctive feature of our approach is the role played by the (not necessarily convex) signal set \mathcal{X} . As indicated in Sect. 3, \mathcal{X} can be viewed as a characteristic of the observation model that allows us to study any situation between non-uniform and (fully) uniform recovery. Note that this refinement is also a crucial ingredient of Proposition 1 and Theorem 3. Nevertheless, it is important to bear in mind that the Gaussian mean width is not trivial to estimate in specific situations, e.g., see (3.3) and Example 1.
- (4) *(Non-)convexity and data-driven priors.* Remarkably, Fact 1 is the only argument in our proof that relies on the convexity of the constraint set K . In principle, it is possible to drop this assumption on K by investigating the projected gradient descent method as an algorithmic implementation of $(P_{K,y})$, e.g., see [46, 47, 56]. However, the feasibility of this approach depends on the existence of an efficient projection onto K .¹² A modern line of research on signal processing with non-convex optimization advocates the use of data-driven priors—a natural consequence of many recent advances on *generative models* in machine learning research, e.g., see [5, 39] and the references therein. Although the algorithmic strategy behind these methods bears resemblance to $(P_{K,y})$, we believe that the complexity of learned signal priors leads to a more difficult mathematical problem, whose understanding is still in its infancy. This is particularly underpinned by the fact that even the situation of convex priors is not sufficiently well understood yet (see [26, 27, 40]).

¹² The existence of an efficient projection and the (non-)convexity of K are not equivalent. There are many examples of efficient projections onto non-convex sets, while the projection onto convex sets can be NP-hard, e.g., see [53].

7 Proof of the Main Result (Theorem 2)

Recall that according to Fact 1, it suffices to show that for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathbb{R}^m$ satisfying $\rho(\mathbf{x}) \leq \frac{t}{32}$ and (2.4), we have that

$$\inf_{\mathbf{v} \in K_{\mathbf{x},t}} \mathcal{E}_{\mathbf{x}}(\mathbf{v}) > 0.$$

To this end, we make use of the decomposition in (2.2), i.e.,

$$\mathcal{E}_{\mathbf{x}}(\mathbf{v}) = \mathcal{Q}(\mathbf{v}) + \mathcal{N}_{\mathbf{x}}(\mathbf{v}) + \mathcal{M}_{\mathbf{x}}(\mathbf{v}),$$

and continue by showing separate bounds for all three terms, where each bound holds uniformly for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{v} \in K_{\mathbf{x},t}$. Unless stated otherwise, we assume that the hypotheses of Theorem 2 are fulfilled throughout this section and we have that $t > 0$; the case of $t = 0$ is in fact much simpler, see Step 4b. Moreover, we will use the notation $K_{\mathcal{X},t} := \cup_{\mathbf{x} \in \mathcal{X}} K_{\mathbf{x},t} = (K - T\mathcal{X}) \cap t\mathbb{S}^{p-1}$.

Step 1: Bounding the quadratic term. Let $\mathbf{A} \in \mathbb{R}^{m \times p}$ denote the matrix with row vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$. In order to control the random variable

$$\mathcal{Q}(\mathbf{v}) = \frac{1}{m} \sum_{i=1}^m \langle \mathbf{a}_i, \mathbf{v} \rangle^2 = \left\| \frac{1}{\sqrt{m}} \mathbf{A}\mathbf{v} \right\|_2^2$$

uniformly for all $\mathbf{v} \in K_{\mathcal{X},t}$, we make use of the following recent matrix deviation inequality for sub-Gaussian matrices.

Theorem 6 ([35, Cor. 1.2]) *There exists a universal constant $C_{\mathcal{Q}} > 0$ for which the following holds.*

Let $H \subset \mathbb{R}^p$, and let $\mathbf{A} \in \mathbb{R}^{m \times p}$ be a random matrix with independent, isotropic, sub-Gaussian row vectors $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$. Moreover, assume that $\max_{i \in [m]} \|\mathbf{a}_i\|_{\psi_2} \leq L$ for some $L > 0$. Then for every $u > 0$, we have the following with probability at least $1 - 3 \exp(-u^2)$:

$$\sup_{\mathbf{v} \in H} \left| \left\| \frac{1}{\sqrt{m}} \mathbf{A}\mathbf{v} \right\|_2 - \|\mathbf{v}\|_2 \right| \leq C_{\mathcal{Q}} \cdot L \sqrt{\log L} \cdot \frac{w(H) + u \cdot \sup_{\mathbf{v} \in H} \|\mathbf{v}\|_2}{\sqrt{m}}.$$

We now apply this result to $H := K_{\mathcal{X},t} \subset t\mathbb{S}^{p-1}$ within the setting of Theorem 2: According to the first branch of the assumption (2.3), we have that $m \geq C \cdot L^2 \log L \cdot (w_t^2(K - T\mathcal{X}) + u^2)$. Hence, by adjusting the universal constant C (only depending on $C_{\mathcal{Q}}$), Theorem 6 implies that the following holds with probability at least $1 - 3 \exp(-u^2)$:

$$\frac{3t^2}{2} \geq \mathcal{Q}(\mathbf{v}) \geq \frac{t^2}{2} \quad \text{for all } \mathbf{v} \in K_{\mathcal{X},t}. \tag{7.1}$$

Step 2: Bounding the noise term. In order to control the noise term $\mathcal{N}_{\mathbf{x}}(\mathbf{v})$, we require the following uniform upper bound for subsums of the quadratic term.

Theorem 7 ([17, Thm. 2.10]) *There exist universal constants $c, C_{\mathcal{N}} > 0$ for which the following holds.*

Let $H \subset \mathbb{R}^p$, and let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ be independent, isotropic, sub-Gaussian random vectors such that $\max_{i \in [m]} \|\mathbf{a}_i\|_{\psi_2} \leq L$ for some $L > 0$. Then for every $m_0 \in \{0, 1, \dots, m\}$ and $u_0 \geq \sqrt{m_0 \log(em/m_0)}$, we have the following with probability at least $1 - 2 \exp(-cu_0^2)$:

$$\sup_{\mathbf{v} \in H} \max_{\substack{\mathcal{I} \subset [m] \\ |\mathcal{I}| \leq m_0}} \left(\frac{1}{m} \sum_{i \in \mathcal{I}} |\langle \mathbf{a}_i, \mathbf{v} \rangle|^2 \right)^{1/2} \leq C_{\mathcal{N}} \cdot L^2 \cdot \frac{w(H) + u_0 \cdot \sup_{\mathbf{v} \in H} \|\mathbf{v}\|_2}{\sqrt{m}}.$$

For $\mathbf{w} \in \mathbb{R}^m$, we denote by $\mathcal{I}_{m_0}(\mathbf{w}) \subset [m]$ any (possibly non-unique) index set that corresponds to the m_0 largest entries of \mathbf{w} in magnitude, i.e., for all $i \in \mathcal{I}_{m_0}(\mathbf{w})$ and $i' \in \mathcal{I}_{m_0}(\mathbf{w})^c$, we have that $|w_i| \geq |w_{i'}|$; note that for $m_0 = 0$, we simply have $\mathcal{I}_{m_0}(\mathbf{w}) = \emptyset$. With this notation at hand, observe that

$$\|\mathbf{w}\|_{[m_0]} = \left(\sum_{i \in \mathcal{I}_{m_0}(\mathbf{w})} |w_i|^2 \right)^{1/2} \quad \text{and} \quad \sigma_{m_0}(\mathbf{w})_2 = \left(\sum_{i \in \mathcal{I}_{m_0}(\mathbf{w})^c} |w_i|^2 \right)^{1/2}.$$

Now, consider the specific choice $\mathbf{w} := \mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})$. The Cauchy–Schwarz inequality then implies

$$\begin{aligned} \sup_{\mathbf{v} \in K_{\mathbf{x},t}} |\mathcal{N}_{\mathbf{x}}(\mathbf{v})| &\leq \frac{2}{m} \sup_{\mathbf{v} \in K_{\mathbf{x},t}} \sum_{i \in \mathcal{I}_{m_0}(\mathbf{w})} |w_i| \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle| \\ &\quad + \frac{2}{m} \sup_{\mathbf{v} \in K_{\mathbf{x},t}} \sum_{i \in \mathcal{I}_{m_0}(\mathbf{w})^c} |w_i| \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle| \\ &\leq \frac{2}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_{[m_0]} \cdot \sup_{\mathbf{v} \in K_{\mathcal{X},t}} \max_{\substack{\mathcal{I} \subset [m] \\ |\mathcal{I}| \leq m_0}} \left(\frac{1}{m} \sum_{i \in \mathcal{I}} |\langle \mathbf{a}_i, \mathbf{v} \rangle|^2 \right)^{1/2} \\ &\quad + \frac{2}{\sqrt{m}} \sigma_{m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}))_2 \cdot \sup_{\mathbf{v} \in K_{\mathcal{X},t}} \sqrt{Q(\mathbf{v})}. \end{aligned}$$

Let us now estimate the first summand of this bound. According to Theorem 7, there exist universal constants $c, C_{\mathcal{N}} > 0$ such that for every $u_0 \geq \sqrt{m_0 \log(em/m_0)}$, we have the following with probability at least $1 - 2 \exp(-cu_0^2)$:

$$\begin{aligned} \sup_{\mathbf{v} \in K_{\mathcal{X},t}} \max_{\substack{\mathcal{I} \subset [m] \\ |\mathcal{I}| \leq m_0}} \left(\frac{1}{m} \sum_{i \in \mathcal{I}} |\langle \mathbf{a}_i, \mathbf{v} \rangle|^2 \right)^{1/2} &\leq C_{\mathcal{N}} \cdot L^2 \cdot \frac{w(K_{\mathcal{X},t}) + u_0 \cdot t}{\sqrt{m}} \\ &\leq \frac{1}{16} \Delta^{-1} t, \end{aligned}$$

where the last inequality follows from the second branch of (2.3). Hence, taking a union bound with the event of (7.1), it follows that with probability at least $1 - 3 \exp(-u^2) - 2 \exp(-cu_0^2)$, the following holds uniformly for every $\mathbf{x} \in \mathcal{X}$ and every

$\mathbf{y} \in \mathbb{R}^m$ satisfying the condition (2.4):

$$\sup_{\mathbf{v} \in K_{\mathbf{x},t}} |\mathcal{M}_{\mathbf{x}}(\mathbf{v})| \leq 2 \cdot \frac{t}{16} \cdot t + 2 \cdot \sqrt{\frac{3}{2}} \cdot t \cdot \frac{t}{20} \leq \frac{t^2}{4}. \tag{7.2}$$

Step 3: Bounding the multiplier term. Our goal is to bound the random variable

$$\frac{1}{2} \mathcal{M}_{\mathbf{x}}(\mathbf{v}) = \frac{1}{m} \sum_{i=1}^m (\langle \mathbf{a}_i, T\mathbf{x} \rangle - \tilde{y}_i(\mathbf{x})) \langle \mathbf{a}_i, \mathbf{v} \rangle$$

uniformly from below for all $\mathbf{v} \in K_{\mathbf{x},t}$ and $\mathbf{x} \in \mathcal{X}$ that satisfy $\rho(\mathbf{x}) \leq \frac{t}{32}$. Let us begin by adding and subtracting the expected value:

$$\frac{1}{2} \mathcal{M}_{\mathbf{x}}(\mathbf{v}) \geq \left(\inf_{\mathbf{v} \in K_{\mathbf{x},t}} \frac{1}{2} \mathbb{E}[\mathcal{M}_{\mathbf{x}}(\mathbf{v})] \right) - \frac{1}{2} |\mathcal{M}_{\mathbf{x}}(\mathbf{v}) - \mathbb{E}[\mathcal{M}_{\mathbf{x}}(\mathbf{v})]|. \tag{7.3}$$

Since \mathbf{a} is isotropic, we observe that $\frac{1}{2} \mathbb{E}[\mathcal{M}_{\mathbf{x}}(\mathbf{v})] = \langle T\mathbf{x} - \mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}], \mathbf{v} \rangle$. Combining this with Definition 1, we obtain a lower bound in terms of the target mismatch:

$$\begin{aligned} \inf_{\mathbf{v} \in K_{\mathbf{x},t}} \frac{1}{2} \mathbb{E}[\mathcal{M}_{\mathbf{x}}(\mathbf{v})] &\geq t \cdot \inf_{\mathbf{v} \in \mathbb{S}^{p-1}} \langle T\mathbf{x} - \mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}], \mathbf{v} \rangle \\ &= -t \cdot \sup_{\mathbf{v} \in \mathbb{S}^{p-1}} \langle \mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] - T\mathbf{x}, \mathbf{v} \rangle \\ &= -t \cdot \|\mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] - T\mathbf{x}\|_2 = -t \cdot \rho(\mathbf{x}). \end{aligned} \tag{7.4}$$

We now turn to the centered multiplier term in (7.3). To this end, recall Assumption 2 and let $\{\tilde{y}_{t,i}(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ be independent copies of the class $\{\tilde{y}_t(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ for $i = 1, \dots, m$.¹³ Inserting these observation variables and applying the triangle inequality then yield

$$\begin{aligned} &\frac{1}{2} |\mathcal{M}_{\mathbf{x}}(\mathbf{v}) - \mathbb{E}[\mathcal{M}_{\mathbf{x}}(\mathbf{v})]| \\ &\leq \frac{1}{m} \left| \sum_{i=1}^m \underbrace{(\langle \mathbf{a}_i, T\mathbf{x} \rangle - \tilde{y}_{t,i}(\mathbf{x})) \langle \mathbf{a}_i, \mathbf{v} \rangle}_{=: \xi_{t,i}(\mathbf{x})} - \mathbb{E}[(\langle \mathbf{a}_i, T\mathbf{x} \rangle - \tilde{y}_{t,i}(\mathbf{x})) \langle \mathbf{a}_i, \mathbf{v} \rangle] \right| \\ &\quad + \mathbb{E}[|\underbrace{\tilde{y}_t(\mathbf{x}) - \tilde{y}_t(\mathbf{x})}_{=: \varepsilon_t(\mathbf{x})}| \cdot |\langle \mathbf{a}, \mathbf{v} \rangle|] + \frac{1}{m} \sum_{i=1}^m |\underbrace{\tilde{y}_{t,i}(\mathbf{x}) - \tilde{y}_t(\mathbf{x})}_{=: \varepsilon_{t,i}(\mathbf{x})}| \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle|. \end{aligned} \tag{7.5}$$

Note that the second empirical product process in (7.5) is not centered yet due to pulling in the absolute values—a crucial step to ensure that the resulting factors have sub-Gaussian increments. Adding and subtracting the expected value of the last summand

¹³ To be more precise, we assume that the tuples $(\mathbf{a}_i, F_i, \tilde{y}_{t,i})$ are independent copies of $(\mathbf{a}, F, \tilde{y}_t)$ for $i = 1, \dots, m$. In particular, the respective conditions of Assumption 2(a)–(c) are also satisfied for $\{\tilde{y}_{t,i}(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$.

in (7.5) and using Assumption 2(a), we obtain the following upper bound:

$$\begin{aligned} \frac{1}{2} |\mathcal{M}_x(\mathbf{v}) - \mathbb{E}[\mathcal{M}_x(\mathbf{v})]| &\leq \frac{1}{m} \left| \sum_{i=1}^m \xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle - \mathbb{E}[\xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle] \right| \\ &+ \frac{1}{m} \left| \sum_{i=1}^m \varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle| - \mathbb{E}[\varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle|] \right| + \underbrace{2\mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |\langle \mathbf{a}, \mathbf{v} \rangle|]}_{\leq \frac{t^2}{32}}. \end{aligned}$$

This estimate in conjunction with (7.4) implies that the following holds uniformly for all $\mathbf{x} \in \mathcal{X}$ with $\rho(\mathbf{x}) \leq \frac{t}{32}$:

$$\begin{aligned} \inf_{\mathbf{v} \in K_{x,t}} \mathcal{M}_x(\mathbf{v}) &\geq -\frac{t^2}{8} - 2 \sup_{\mathbf{v} \in K_{x,t}} \frac{1}{m} \left| \sum_{i=1}^m \xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle - \mathbb{E}[\xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle] \right| \\ &- 2 \sup_{\mathbf{v} \in K_{x,t}} \frac{1}{m} \left| \sum_{i=1}^m \varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle| - \mathbb{E}[\varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle|] \right| \\ &\geq -\frac{t^2}{8} - 2 \sup_{\substack{\mathbf{x} \in \mathcal{X} \\ \mathbf{v} \in K_{\mathcal{X},t}}} \frac{1}{m} \left| \sum_{i=1}^m \xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle - \mathbb{E}[\xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle] \right| \\ &- 2 \sup_{\substack{\mathbf{x} \in \mathcal{X} \\ \mathbf{v} \in K_{\mathcal{X},t}}} \frac{1}{m} \left| \sum_{i=1}^m \varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle| - \mathbb{E}[\varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle|] \right|. \quad (7.6) \end{aligned}$$

By Assumption 2(b) and (c), all factors of the product processes

$$\left\{ \frac{1}{m} \sum_{i=1}^m \xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle \right\}_{\mathbf{x} \in \mathcal{X}, \mathbf{v} \in K_{\mathcal{X},t}}, \quad \left\{ \frac{1}{m} \sum_{i=1}^m \varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle| \right\}_{\mathbf{x} \in \mathcal{X}, \mathbf{v} \in K_{\mathcal{X},t}} \quad (7.7)$$

have sub-Gaussian increments. A key ingredient for controlling these empirical processes is a powerful concentration inequality due to Mendelson [42]. The following result is adapted from [42, Thm. 1.13].

Theorem 8 *There exist universal constants $c, C_{\mathcal{M}} > 0$ for which the following holds.*

Let $\{g_a\}_{a \in \mathcal{A}}$ and $\{h_b\}_{b \in \mathcal{B}}$ be stochastic processes indexed by two sets \mathcal{A} and \mathcal{B} , both defined on a common probability space $(\Omega, \mathcal{A}, \mathbb{P})$. We assume that there exist $r_{\mathcal{A}}, r_{\mathcal{B}} \geq 0$ and pseudo-metrics $d_{\mathcal{A}}$ on \mathcal{A} and $d_{\mathcal{B}}$ on \mathcal{B} such that

$$\begin{aligned} \|g_a - g_{a'}\|_{\psi_2} &\leq d_{\mathcal{A}}(a, a') \quad \text{and} \quad \|g_a\|_{\psi_2} \leq r_{\mathcal{A}} \quad \text{for all } a, a' \in \mathcal{A}, \\ \|h_b - h_{b'}\|_{\psi_2} &\leq d_{\mathcal{B}}(b, b') \quad \text{and} \quad \|h_b\|_{\psi_2} \leq r_{\mathcal{B}} \quad \text{for all } b, b' \in \mathcal{B}. \end{aligned}$$

Finally, let X_1, \dots, X_m be independent copies of a random variable $X \sim \mathbb{P}$. Then for every $u \geq 1$, we have the following with probability at least $1 - 2 \exp(-cu^2)$:

$$\begin{aligned} & \sup_{\substack{a \in \mathcal{A} \\ b \in \mathcal{B}}} \frac{1}{m} \left| \sum_{i=1}^m g_a(X_i) h_b(X_i) - \mathbb{E}[g_a(X_i) h_b(X_i)] \right| \\ & \leq C_{\mathcal{M}} \cdot \left(\frac{(\gamma_2(\mathcal{A}, d_{\mathcal{A}}) + u \cdot r_{\mathcal{A}}) \cdot (\gamma_2(\mathcal{B}, d_{\mathcal{B}}) + u \cdot r_{\mathcal{B}})}{m} \right. \\ & \quad \left. + \frac{r_{\mathcal{A}} \cdot \gamma_2(\mathcal{B}, d_{\mathcal{B}}) + r_{\mathcal{B}} \cdot \gamma_2(\mathcal{A}, d_{\mathcal{A}}) + u \cdot r_{\mathcal{A}} r_{\mathcal{B}}}{\sqrt{m}} \right), \end{aligned}$$

where $\gamma_2(\cdot, \cdot)$ denotes Talagrand’s γ_2 -functional (e.g., see [42, Def. 1.2]).

Let us now apply Theorem 8 to the first product process in (7.7). To this end, we observe that the class $\{\langle \mathbf{a}, \mathbf{v} \rangle\}_{\mathbf{v} \in K_{\mathcal{X},t}}$ has sub-Gaussian increments with respect to the pseudo-metric induced by $L \|\cdot\|_2$, while Assumption 2(b) implies that $\{\xi_t(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ has sub-Gaussian increments with respect to $L_t \cdot d_T$. Hence, Theorem 8 implies that the following holds with probability at least $1 - 2 \exp(-cu^2)$:¹⁴

$$\begin{aligned} & \sup_{\substack{\mathbf{x} \in \mathcal{X} \\ \mathbf{v} \in K_{\mathcal{X},t}}} \frac{1}{m} \left| \sum_{i=1}^m \xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle - \mathbb{E}[\xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle] \right| \\ & \leq C_{\mathcal{M}} \cdot \left(\frac{(L_t \gamma_2(\mathcal{X}, d_T) + u \cdot r) \cdot (L \gamma_2(K_{\mathcal{X},t}, \|\cdot\|_2) + u \cdot Lt)}{m} \right. \\ & \quad \left. + \frac{r \cdot L \gamma_2(K_{\mathcal{X},t}, \|\cdot\|_2) + Lt \cdot L_t \gamma_2(\mathcal{X}, d_T) + u \cdot r Lt}{\sqrt{m}} \right). \end{aligned}$$

According to Talagrand’s majorizing measure theorem [59, Thm. 2.4.1], we have that

$$w(H) \asymp \gamma_2(H, \|\cdot\|_2) \quad \text{for every } H \subset \mathbb{R}^p.$$

Moreover, it is not hard to see that $\gamma_2(\mathcal{X}, d_T) = \gamma_2(T\mathcal{X}, \|\cdot\|_2)$. Consequently, the following bound holds with probability at least $1 - 2 \exp(-cu^2)$:

$$\begin{aligned} & \sup_{\substack{\mathbf{x} \in \mathcal{X} \\ \mathbf{v} \in K_{\mathcal{X},t}}} \frac{1}{m} \left| \sum_{i=1}^m \xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle - \mathbb{E}[\xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle] \right| \\ & \leq C'_{\mathcal{M}} \cdot t \cdot L \cdot \left(\frac{(L_t w(T\mathcal{X}) + u \cdot r) \cdot (w_t(K - T\mathcal{X}) + u)}{m} \right. \\ & \quad \left. + \frac{r \cdot w_t(K - T\mathcal{X}) + L_t w(T\mathcal{X}) + u \cdot r}{\sqrt{m}} \right). \end{aligned} \tag{7.8}$$

¹⁴ To be slightly more precise, we apply Theorem 8 to the index sets $\mathcal{A} := \mathcal{X}$ and $\mathcal{B} := K_{\mathcal{X},t}$, while $X \sim \mathbb{P}$ corresponds to the random tuple $(\mathbf{a}, F, \tilde{y}_t)$.

The second product process in (7.7) can be treated similarly: The class $\{|\langle \mathbf{a}, \mathbf{v} \rangle|\}_{\mathbf{v} \in K_{\mathcal{X},t}}$ has sub-Gaussian increments with respect to $L \|\cdot\|_2$, while Assumption 2(c) implies that $\{\varepsilon_t(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ has sub-Gaussian increments with respect to $\hat{L}_t \cdot d_T$. An analogous application of Theorem 8 shows that with probability at least $1 - 2 \exp(-cu^2)$, we have that

$$\begin{aligned} & \sup_{\substack{\mathbf{x} \in \mathcal{X} \\ \mathbf{v} \in K_{\mathcal{X},t}}} \frac{1}{m} \left| \sum_{i=1}^m \varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle| - \mathbb{E}[\varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle|] \right| \\ & \leq C'_{\mathcal{M}} \cdot t \cdot L \cdot \left(\frac{(\hat{L}_t w(T\mathcal{X}) + u \cdot \hat{r}) \cdot (w_t(K - T\mathcal{X}) + u)}{m} \right. \\ & \quad \left. + \frac{\hat{r} \cdot w_t(K - T\mathcal{X}) + \hat{L}_t w(T\mathcal{X}) + u \cdot \hat{r}}{\sqrt{m}} \right). \end{aligned} \tag{7.9}$$

We are ready to prove our main result.

Step 4a: Proof of Theorem 2 ($t > 0$). We now assume that the events corresponding to (7.1), (7.2), (7.8), and (7.9) have occurred jointly with probability at least $1 - 7 \exp(-cu^2) - 2 \exp(-cu_0^2)$ for an appropriate constant $c > 0$; note that the factors 7 and 2 can be removed by slightly adjusting c . Combining these bounds with (7.6), the following holds uniformly for every $\mathbf{x} \in \mathcal{X}$ with $\rho(\mathbf{x}) \leq \frac{1}{32}$ and every $\mathbf{y} \in \mathbb{R}^m$ satisfying the condition (2.4):

$$\begin{aligned} \inf_{\mathbf{v} \in K_{\mathcal{X},t}} \mathcal{E}_{\mathbf{x}}(\mathbf{v}) & \geq \inf_{\mathbf{v} \in K_{\mathcal{X},t}} \mathcal{Q}(\mathbf{v}) - \sup_{\mathbf{v} \in K_{\mathcal{X},t}} |\mathcal{N}_{\mathbf{x}}(\mathbf{v})| + \inf_{\mathbf{v} \in K_{\mathcal{X},t}} \mathcal{M}_{\mathbf{x}}(\mathbf{v}) \\ & \geq \frac{t^2}{2} - \frac{t^2}{4} - \frac{t^2}{8} \\ & \quad - C'_{\mathcal{M}} \cdot t \cdot L \cdot \left(\frac{(r + \hat{r}) \cdot w_t(K - T\mathcal{X}) + (L_t + \hat{L}_t) \cdot w(T\mathcal{X}) + u \cdot (r + \hat{r})}{\sqrt{m}} \right) \\ & \quad - C'_{\mathcal{M}} \cdot t \cdot L \cdot \left(\frac{((L_t + \hat{L}_t) \cdot w(T\mathcal{X}) + u \cdot (r + \hat{r})) \cdot (w_t(K - T\mathcal{X}) + u)}{m} \right). \end{aligned}$$

If we could show that this lower bound is strictly positive, the claim of Theorem 2 would follow directly from Fact 1. To conclude this argument, it is enough to have that

$$m \geq C' \cdot L^2 t^{-2} \cdot \left((r + \hat{r}) \cdot w_t(K - T\mathcal{X}) + (L_t + \hat{L}_t) \cdot w(T\mathcal{X}) + u \cdot (r + \hat{r}) \right)^2, \tag{7.10}$$

$$m \geq C' \cdot L t^{-1} \cdot ((L_t + \hat{L}_t) \cdot w(T\mathcal{X}) + u \cdot (r + \hat{r})) \cdot (w_t(K - T\mathcal{X}) + u), \tag{7.11}$$

where $C' > 0$ is an appropriate universal constant. Indeed, both (7.10) and (7.11) are consequences of (2.3): The bound of (7.10) is equivalent to the third branch of (2.3), while (7.11) follows from the multiplication of the first and third branch of (2.3) and then taking the square root. Note that the previous argument also makes use of the basic fact that $(v + w)^2 \geq v^2 + w^2 \geq \frac{1}{2}(v + w)^2$ for all $v, w \geq 0$, and that $L^2 \gtrsim L$ due to the isotropy of \mathbf{a} .

Step 4b: Proof of Theorem 2 ($t = 0$). We may assume that $r = \hat{r} = L_t = \hat{L}_t = 0$, $m_0 = 0$, $\Delta = L^{-1}\sqrt{\log L}$, and $u = u_0$, while only considering those $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathbb{R}^m$ with $\mathbf{y} = \tilde{\mathbf{y}}(\mathbf{x})$ and $\rho(\mathbf{x}) = 0$. Consequently, we have a noiseless linear model, i.e., $\mathbf{y} = \tilde{\mathbf{y}}(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle$. It follows that $\mathcal{N}_{\mathbf{x}}(\cdot) = \mathcal{M}_{\mathbf{x}}(\cdot) = 0$. Analogously to (7.1), we can conclude that with probability at least $1 - 3 \exp(-u^2)$, the quadratic term satisfies $\mathcal{Q}(\mathbf{v}) \geq \frac{1}{2}$ for all $\mathbf{v} \in K_{\mathcal{X},0} := \text{cone}(K - T\mathcal{X}) \cap \mathbb{S}^{p-1}$. On this event, let $\hat{\mathbf{z}} \in K$ be any minimizer of $(P_{K,\mathbf{y}})$ with input $\mathbf{y} = \tilde{\mathbf{y}}(\mathbf{x})$ for some $\mathbf{x} \in \mathcal{X}$. If we would not have exact recovery, i.e., $\hat{\mathbf{v}} := \hat{\mathbf{z}} - T\mathbf{x} \neq \mathbf{0}$, then

$$\mathcal{E}_{\mathbf{x}}(\hat{\mathbf{v}}) = \mathcal{Q}(\hat{\mathbf{v}}) = \|\hat{\mathbf{v}}\|_2^2 \cdot \mathcal{Q}\left(\underbrace{\frac{\hat{\mathbf{v}}}{\|\hat{\mathbf{v}}\|_2}}_{\in K_{\mathcal{X},0}}\right) \geq \frac{\|\hat{\mathbf{v}}\|_2^2}{2} > 0, \tag{7.12}$$

which contradicts the fact that $\hat{\mathbf{z}}$ is a solution to $(P_{K,\mathbf{y}})$. □

8 Proofs for Sect. 4

Each result in Sect. 4 is an application of Theorem 2 and follows the same proof template:

Step 1. How the prerequisites of Theorem 2 are met:

Step 1a. Implementation of Assumption 1.

Step 1b. Controlling the target mismatch $\rho(\mathbf{x})$ for every $\mathbf{x} \in \mathcal{X}$.

Step 1c. Controlling the increment parameters of Assumption 2.

Step 2. Proof of the actual statement via Theorem 2.

Step 3. Verification of Step 1b.

Step 4. Verification of Step 1c.

8.1 Proofs for Sect. 4.1

Proof of Corollary 2 We follow the proof template from the beginning of Sect. 8:

Step 1a. The model setup of Corollary 2 fits into Assumption 1 as follows:

- We have that $\mathbf{a} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ and therefore $\|\mathbf{a}\|_{\psi_2} \lesssim 1$. The signal set \mathcal{X} is an arbitrary subset of \mathbb{R}^p . The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(\mathbf{a}, \mathbf{x}) := \text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle)$.

- The target function $T: \mathcal{X} \rightarrow K$ corresponds to the (scaled) normalization

$$T\mathbf{x} := \sqrt{\frac{2}{\pi}} \frac{\mathbf{x}}{\|\mathbf{x}\|_2}. \text{ In particular, we have that } d_T(\mathbf{x}, \mathbf{x}') = \sqrt{\frac{2}{\pi}} \left\| \frac{\mathbf{x}}{\|\mathbf{x}\|_2} - \frac{\mathbf{x}'}{\|\mathbf{x}'\|_2} \right\|_2.$$

Step 1b. The target mismatch $\rho(\mathbf{x})$ vanishes for every $\mathbf{x} \in \mathcal{X}$.

Step 1c. There exists an approximation $\tilde{\mathbf{y}}_t(\mathbf{x})$ of the observation variable $\tilde{\mathbf{y}}(\mathbf{x})$ such that the conditions of Assumption 2 are fulfilled with

$$L_t \lesssim 1 + t^{-1}, \quad \hat{L}_t \lesssim t^{-1}, \quad r \lesssim 1, \quad \hat{r} \lesssim 1. \tag{8.1}$$

Step 2. The first and third branch of the condition (2.3) is implied by (4.2) for a sufficiently large universal constant $C' > 0$. Since $\phi(\beta) := \beta\sqrt{\log(e/\beta)}$ defines a continuous and non-decreasing function on $[0, 1]$ with $\phi(1) = 1$ and $\phi(0) := 0$ (by continuous extension), we may assume without loss of generality that $\beta\sqrt{\log(e/\beta)} = c_0 t \in (0, 1]$. Now, we set

$$\Delta^2 := \frac{1}{t\sqrt{\log(e/\beta)}}, \quad m_0 := \lfloor \beta m \rfloor, \quad u_0 := \sqrt{2m\beta \log(e/\beta)},$$

implying that $u_0 \geq \sqrt{m_0 \log(em/m_0)}$ and $u_0^2 \geq 2c_0 t m$. Combining the latter inequality with (4.2) for C' sufficiently large particularly implies that $u_0^2 \geq u^2$. Furthermore, observe that $\Delta^2 u_0^2 = 2c_0 m$ and $\Delta^2 \cdot w_t^2(K - T\mathcal{X}) \leq t^{-1} \cdot w_t^2(K - T\mathcal{X})$. Hence, the second branch of the condition (2.3) is satisfied if c_0 is chosen small enough and C' in (4.2) large enough.

Next, we show that every $\mathbf{y} \in \{-1, 1\}^m$ given by (4.1) with $\frac{1}{2m} \|\mathbf{v}\|_1 \leq \beta$ also satisfies (2.4). Since $\mathbf{v} = \mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}) \in \{-2, 0, 2\}^m$ and $\frac{1}{2} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_1 \leq \beta m$, it follows that $\{|i \in [m] : y_i \neq \tilde{y}_i(\mathbf{x})\} \leq \lfloor \beta m \rfloor$. Consequently, $\sigma_{m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}))_2 = 0$ and

$$\frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_{[m_0]} = \frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_2 \leq 2\sqrt{\beta} = 2\left(\frac{c_0 t}{\sqrt{\log(e/\beta)}}\right)^{1/2} \leq \Delta t,$$

where the last inequality holds for $c_0 \leq \frac{1}{4}$. Theorem 2 now implies the claim of Corollary 2.

Step 3. Let $\mathbf{x} \in \mathcal{X}$ and consider the orthogonal decomposition of the standard Gaussian random vector \mathbf{a} along \mathbf{x} :

$$\mathbf{a} = \langle \mathbf{a}, \bar{\mathbf{x}} \rangle \bar{\mathbf{x}} + P_{\mathbf{x}^\perp}(\mathbf{a}),$$

where $\bar{\mathbf{x}} := \frac{\mathbf{x}}{\|\mathbf{x}\|_2}$ and $P_{\mathbf{x}^\perp} := P_{\{\mathbf{x}\}^\perp}$. Since $\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle)$ is centered and $\langle \mathbf{a}, \bar{\mathbf{x}} \rangle \sim N(0, 1)$ is independent of $P_{\mathbf{x}^\perp}(\mathbf{a})$, we have that

$$\begin{aligned} \mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] &= \mathbb{E}[\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle)(\langle \mathbf{a}, \bar{\mathbf{x}} \rangle \bar{\mathbf{x}} + P_{\mathbf{x}^\perp}(\mathbf{a}))] \\ &= \mathbb{E}[|\langle \mathbf{a}, \bar{\mathbf{x}} \rangle|] \cdot \bar{\mathbf{x}} + \mathbb{E}[\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle)] \cdot \mathbb{E}[P_{\mathbf{x}^\perp}(\mathbf{a})] = \sqrt{\frac{2}{\pi}} \cdot \bar{\mathbf{x}} = T\mathbf{x}, \end{aligned}$$

which implies that $\rho(\mathbf{x}) = 0$.

Step 4. Using the shortcut $\bar{\mathbf{x}} := \frac{\mathbf{x}}{\|\mathbf{x}\|_2}$ again, we approximate the observation variable $\tilde{\mathbf{y}}(\mathbf{x}) = \text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle)$ by $\tilde{\mathbf{y}}_t(\mathbf{x}) := \psi_t(\langle \mathbf{a}, \bar{\mathbf{x}} \rangle)$ for $t \in (0, 1]$ and $\mathbf{x} \in \mathcal{X}$ where

$$\psi_t(s) := \begin{cases} \frac{128}{t} \cdot s, & |s| \leq \frac{t}{128}, \\ \text{sign}(s), & \text{otherwise,} \end{cases} \quad s \in \mathbb{R}.$$

See also Fig. 2 for an illustration. Since $\tilde{y}(\mathbf{x}) = \text{sign}(\langle \mathbf{a}, \bar{\mathbf{x}} \rangle)$, the absolute value of the approximation error then takes the form

$$\varepsilon_t(\mathbf{x}) = |\text{sign}(\langle \mathbf{a}, \bar{\mathbf{x}} \rangle) - \psi_t(\langle \mathbf{a}, \bar{\mathbf{x}} \rangle)| \cdot \chi_{[0, t^\circ]}(|\langle \mathbf{a}, \bar{\mathbf{x}} \rangle|),$$

where we have used $t^\circ := \frac{t}{128}$ for the sake of notational convenience. We now show that for this choice of approximation, the conditions of Assumption 2 are indeed fulfilled with (8.1):

On Assumption 2(a). Let $\mathbf{x} \in \mathcal{X}$, $\mathbf{z} \in \mathbb{S}^{p-1}$, and consider the orthogonal decomposition

$$\mathbf{z} = \langle \mathbf{z}, \bar{\mathbf{x}} \rangle \bar{\mathbf{x}} + P_{\bar{\mathbf{x}}^\perp}(\mathbf{z}).$$

This implies

$$\begin{aligned} \mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] &\leq \mathbb{E}[\chi_{[0, t^\circ]}(|\langle \mathbf{a}, \bar{\mathbf{x}} \rangle|) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] \\ &\leq |\langle \mathbf{z}, \bar{\mathbf{x}} \rangle| \cdot \mathbb{E}[\chi_{[0, t^\circ]}(|\langle \mathbf{a}, \bar{\mathbf{x}} \rangle|) \cdot |\langle \mathbf{a}, \bar{\mathbf{x}} \rangle|] \\ &\quad + \mathbb{E}[\chi_{[0, t^\circ]}(|\langle \mathbf{a}, \bar{\mathbf{x}} \rangle|) \cdot |\langle \mathbf{a}, P_{\bar{\mathbf{x}}^\perp}(\mathbf{z}) \rangle|]. \end{aligned}$$

Clearly, $|\langle \mathbf{z}, \bar{\mathbf{x}} \rangle| \leq 1$ and $\mathbb{E}[\chi_{[0, t^\circ]}(|\langle \mathbf{a}, \bar{\mathbf{x}} \rangle|) \cdot |\langle \mathbf{a}, \bar{\mathbf{x}} \rangle|] \leq t^\circ$. Since the random variables $\langle \mathbf{a}, \bar{\mathbf{x}} \rangle$ and $\langle \mathbf{a}, P_{\bar{\mathbf{x}}^\perp}(\mathbf{z}) \rangle$ are independent, and $\langle \mathbf{a}, \bar{\mathbf{x}} \rangle$ is standard Gaussian, we obtain

$$\begin{aligned} \mathbb{E}[\chi_{[0, t^\circ]}(|\langle \mathbf{a}, \bar{\mathbf{x}} \rangle|) \cdot |\langle \mathbf{a}, P_{\bar{\mathbf{x}}^\perp}(\mathbf{z}) \rangle|] &= \mathbb{P}(|\langle \mathbf{a}, \bar{\mathbf{x}} \rangle| \leq t^\circ) \cdot \mathbb{E}[|\langle \mathbf{a}, P_{\bar{\mathbf{x}}^\perp}(\mathbf{z}) \rangle|] \\ &\leq t^\circ \cdot \mathbb{E}[|\langle \mathbf{a}, P_{\bar{\mathbf{x}}^\perp}(\mathbf{z}) \rangle|]. \end{aligned}$$

Moreover, by Jensen’s inequality and the isotropy of \mathbf{a} , it follows that

$$\mathbb{E}[|\langle \mathbf{a}, P_{\bar{\mathbf{x}}^\perp}(\mathbf{z}) \rangle|] \leq (\mathbb{E}[|\langle \mathbf{a}, P_{\bar{\mathbf{x}}^\perp}(\mathbf{z}) \rangle|^2])^{1/2} = \|P_{\bar{\mathbf{x}}^\perp}(\mathbf{z})\|_2 \leq 1.$$

Putting everything together, this shows that Assumption 2(a) is satisfied.

On Assumption 2(b). Since ψ_t is $\frac{128}{t}$ -Lipschitz, the following holds for every $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ (with $\bar{\mathbf{x}}' := \frac{\mathbf{x}'}{\|\mathbf{x}'\|_2}$):

$$\begin{aligned} \|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} &\leq \sqrt{\frac{2}{\pi}} \|\langle \mathbf{a}, \bar{\mathbf{x}} - \bar{\mathbf{x}}' \rangle\|_{\psi_2} + \|\psi_t(\langle \mathbf{a}, \bar{\mathbf{x}} \rangle) - \psi_t(\langle \mathbf{a}, \bar{\mathbf{x}}' \rangle)\|_{\psi_2} \\ &\leq \left(\sqrt{\frac{2}{\pi}} + \frac{128}{t}\right) \cdot \|\langle \mathbf{a}, \bar{\mathbf{x}} - \bar{\mathbf{x}}' \rangle\|_{\psi_2} \lesssim (1 + t^{-1}) \cdot d_T(\mathbf{x}, \mathbf{x}'). \end{aligned}$$

This implies $L_t \lesssim 1 + t^{-1}$. Furthermore, observe that $|\tilde{y}_t(\mathbf{x})| \leq 1$, so that

$$\|\xi_t(\mathbf{x})\|_{\psi_2} \leq \sqrt{\frac{2}{\pi}} \|\langle \mathbf{a}, \bar{\mathbf{x}} \rangle\|_{\psi_2} + \|\tilde{y}_t(\mathbf{x})\|_{\psi_2} \lesssim 1$$

for every $\mathbf{x} \in \mathcal{X}$. This shows $r \lesssim 1$.

On Assumption 2(c). We observe that the function $s \mapsto \phi_t(s) := |\text{sign}(s) - \psi_t(s)|$ is $\frac{128}{t}$ -Lipschitz. Therefore, for every $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, we obtain

$$\|\varepsilon_t(\mathbf{x}) - \varepsilon_t(\mathbf{x}')\|_{\psi_2} \leq \frac{128}{t} \cdot \|\langle \mathbf{a}, \bar{\mathbf{x}} - \bar{\mathbf{x}}' \rangle\|_{\psi_2} \lesssim t^{-1} \cdot d_T(\mathbf{x}, \mathbf{x}'),$$

which implies $\hat{L}_t \lesssim t^{-1}$. Finally, since $|\varepsilon_t(\mathbf{x})| \leq 1$ for every $\mathbf{x} \in \mathcal{X}$, it follows that $\hat{r} \lesssim 1$. □

8.2 Proofs for Sect. 4.2

Proof of Corollary 3 We follow the proof template from the beginning of Sect. 8:

Step 1a. The model setup of Corollary 3 fits into Assumption 1 as follows:

- The measurement vector $\mathbf{a} \in \mathbb{R}^p$ is centered, isotropic, and sub-Gaussian with $\|\mathbf{a}\|_{\psi_2} \leq L$. The signal set \mathcal{X} satisfies $\mathcal{X} \subset R\mathbb{B}_2^p$. The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(\mathbf{a}, \mathbf{x}) := \text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)$, where τ is uniformly distributed on $[-\lambda, \lambda]$ and independent of \mathbf{a} . In particular, F is a random function.
- The target function $T: \mathcal{X} \rightarrow K$ corresponds to rescaling by a factor of λ^{-1} , i.e., $T\mathbf{x} := \lambda^{-1}\mathbf{x}$. In particular, we have that $d_T(\mathbf{x}, \mathbf{x}') = \lambda^{-1}\|\mathbf{x} - \mathbf{x}'\|_2$.

Step 1b. There exists an absolute constant $\tilde{C} \geq e$ such that if

$$\lambda \geq \tilde{C} \cdot R \cdot L \cdot \sqrt{\log(e/t)}, \tag{8.2}$$

the target mismatch satisfies $\rho(\mathbf{x}) \leq \frac{t}{32}$ for every $\mathbf{x} \in \mathcal{X}$.

Step 1c. There exists an approximation $\tilde{y}_t(\mathbf{x})$ of the observation variable $\tilde{y}(\mathbf{x})$ such that the conditions of Assumption 2 are fulfilled with

$$L_t \leq L \cdot (1 + 64t^{-1}), \quad \hat{L}_t \leq 64Lt^{-1}, \quad r \lesssim RL\lambda^{-1} + 1, \quad \hat{r} \lesssim 1. \tag{8.3}$$

Step 2. The first and third branch of the condition (2.3) is implied by (4.4) for a sufficiently large universal constant $C' > 0$. Since $\phi(\beta) := \beta\sqrt{\log(e/\beta)}$ defines a continuous and non-decreasing function on $[0, 1]$ with $\phi(1) = 1$ and $\phi(0) := 0$ (by continuous extension), we may assume without loss of generality that $\beta\sqrt{\log(e/\beta)} = c_0L^{-2}t \in (0, 1]$. Now, we set

$$\Delta^2 := \frac{1}{tL^2\sqrt{\log(e/\beta)}}, \quad m_0 := \lfloor \beta m \rfloor, \quad u_0 := \sqrt{2m\beta \log(e/\beta)},$$

implying that $u_0 \geq \sqrt{m_0 \log(em/m_0)}$ and $u_0^2 \geq 2c_0mL^{-2}t$. Combining the latter inequality with (4.4) for C' sufficiently large particularly implies that $u_0^2 \geq u^2$. Furthermore, we may assume that $c_0 \leq \frac{1}{4}C^{-1}$ and that (4.4) holds with $C' \geq 2C$,

where $C > 0$ denotes the universal constant from (2.3). Then, the second branch of the condition (2.3) is satisfied, since

$$CL^4 \Delta^2 u_0^2 = \frac{CL^2}{t\sqrt{\log(e/\beta)}} \cdot 2\beta m \log(e/\beta) = 2c_0 C \cdot m \leq \frac{m}{2}$$

and

$$\begin{aligned} CL^4 \Delta^2 \cdot w_t^2(K - \lambda^{-1} \mathcal{X}) &= \frac{CL^2}{t\sqrt{\log(e/\beta)}} \cdot w_t^2(K - \lambda^{-1} \mathcal{X}) \\ &\leq \frac{1}{2} C' \cdot L^2 t^{-1} \cdot w_t^2(K - \lambda^{-1} \mathcal{X}) \leq \frac{m}{2}. \end{aligned}$$

Next, we show that every $\mathbf{y} \in \{-1, 1\}^m$ given by (4.3) with $\frac{1}{2m} \|\mathbf{v}\|_1 \leq \beta$ also satisfies (2.4). Since $\mathbf{v} = \mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}) \in \{-2, 0, 2\}^m$ and $\frac{1}{2} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_1 \leq \beta m$, it follows that $|\{i \in [m] : y_i \neq \tilde{y}_i(\mathbf{x})\}| \leq \lfloor \beta m \rfloor$. Consequently, $\sigma_{m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}))_2 = 0$ and

$$\frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_{[m_0]} = \frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_2 \leq 2\sqrt{\beta} = 2\left(\frac{c_0 t}{L^2 \sqrt{\log(e/\beta)}}\right)^{1/2} \leq \Delta t,$$

where the last inequality holds for $c_0 \leq \frac{1}{4}$. Theorem 2 now implies the claim of Corollary 3.

Step 3. Since \mathbf{a} is isotropic, we have that

$$\rho(\mathbf{x}) = \|\mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] - \lambda^{-1}\mathbf{x}\|_2 \leq \|\mathbb{E}[(\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \lambda^{-1}\mathbf{x} \rangle)\mathbf{a}]\|_2.$$

Therefore, it suffices to show that the following holds for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathbb{S}^{p-1}$:

$$\mathbb{E}[(\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \lambda^{-1}\mathbf{x} \rangle)\langle \mathbf{a}, \mathbf{z} \rangle] \leq \frac{t}{32}.$$

The following identity explains why adding a uniformly distributed dither $\tau \in [-\lambda, \lambda]$ before quantization is useful:

$$\mathbb{E}_\tau[\text{sign}(s + \tau)] = \lambda^{-1}s \cdot \chi_{[-\lambda, \lambda]}(s) + \text{sign}(s) \cdot \chi_{\mathbb{R} \setminus [-\lambda, \lambda]}(s), \quad s \in \mathbb{R}.$$

In other words, for s small enough, integrating over the dithering variable τ allows us to “smooth out” the discontinuity of the sign function. As τ and \mathbf{a} are independent, we can apply the above identity as follows:

$$\begin{aligned} &\mathbb{E}[(\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \lambda^{-1}\mathbf{x} \rangle)\langle \mathbf{a}, \mathbf{z} \rangle] \\ &= \mathbb{E}_\mathbf{a} \mathbb{E}_\tau[(\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \lambda^{-1}\mathbf{x} \rangle)\langle \mathbf{a}, \mathbf{z} \rangle] \\ &= \mathbb{E}[(\lambda^{-1}\langle \mathbf{a}, \mathbf{x} \rangle \cdot \chi_{[-\lambda, \lambda]}(\langle \mathbf{a}, \mathbf{x} \rangle) \\ &\quad + \text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle) \cdot \chi_{\mathbb{R} \setminus [-\lambda, \lambda]}(\langle \mathbf{a}, \mathbf{x} \rangle) - \lambda^{-1}\langle \mathbf{a}, \mathbf{x} \rangle)\langle \mathbf{a}, \mathbf{z} \rangle] \\ &= \mathbb{E}[(\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle) - \lambda^{-1}\langle \mathbf{a}, \mathbf{x} \rangle) \cdot \chi_{\mathbb{R} \setminus [-\lambda, \lambda]}(\langle \mathbf{a}, \mathbf{x} \rangle) \cdot \langle \mathbf{a}, \mathbf{z} \rangle]. \end{aligned}$$

Using the Cauchy–Schwarz inequality twice, the isotropy of \mathbf{a} , and the triangle inequality, we now obtain

$$\begin{aligned} & \mathbb{E}[(\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle) - \lambda^{-1} \langle \mathbf{a}, \mathbf{x} \rangle) \cdot \chi_{\mathbb{R} \setminus [-\lambda, \lambda]}(\langle \mathbf{a}, \mathbf{x} \rangle) \cdot \langle \mathbf{a}, \mathbf{z} \rangle] \\ & \leq \left(\mathbb{E}[(\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle) - \lambda^{-1} \langle \mathbf{a}, \mathbf{x} \rangle)^2 \cdot \chi_{\mathbb{R} \setminus [-\lambda, \lambda]}(\langle \mathbf{a}, \mathbf{x} \rangle)] \right)^{1/2} \cdot (\mathbb{E}[\langle \mathbf{a}, \mathbf{z} \rangle^2])^{1/2} \\ & = \left(\mathbb{E}[(\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle) - \lambda^{-1} \langle \mathbf{a}, \mathbf{x} \rangle)^2 \cdot \chi_{\mathbb{R} \setminus [-\lambda, \lambda]}(\langle \mathbf{a}, \mathbf{x} \rangle)] \right)^{1/2} \\ & \leq \|\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle) - \lambda^{-1} \langle \mathbf{a}, \mathbf{x} \rangle\|_{L^4} \cdot (\mathbb{P}(|\langle \mathbf{a}, \mathbf{x} \rangle| > \lambda))^{1/4} \\ & \leq (1 + \lambda^{-1} \|\langle \mathbf{a}, \mathbf{x} \rangle\|_{L^4}) \cdot (\mathbb{P}(|\langle \mathbf{a}, \mathbf{x} \rangle| > \lambda))^{1/4}. \end{aligned}$$

Since \mathbf{a} is sub-Gaussian with $\|\mathbf{a}\|_{\psi_2} \leq L$, there exist absolute constants $C'' \geq 1$, $c'' > 0$ such that $\|\langle \mathbf{a}, \mathbf{x} \rangle\|_{L^4} \leq C'' \cdot \|\langle \mathbf{a}, \mathbf{x} \rangle\|_{\psi_2} \leq C'' \cdot L \cdot \|\mathbf{x}\|_2$ and $\mathbb{P}(|\langle \mathbf{a}, \mathbf{x} \rangle| > \lambda) \leq 2 \exp(-c'' \lambda^2 (L \|\mathbf{x}\|_2)^{-2})$. Finally, using that every $\mathbf{x} \in \mathcal{X}$ satisfies $\|\mathbf{x}\|_2 \leq R$, we can conclude that

$$\begin{aligned} \mathbb{E}[(\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \lambda^{-1} \mathbf{x} \rangle) \langle \mathbf{a}, \mathbf{z} \rangle] & \leq C'' \cdot (1 + RL\lambda^{-1}) \cdot \exp(-c'' \lambda^2 (RL)^{-2}) \\ & \leq \frac{t}{32}, \end{aligned}$$

where the last estimate is due to (8.2) for $\tilde{C} \geq e$ large enough.

Step 4. We approximate the observation variable $\tilde{y}(\mathbf{x}) = \text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)$ by $\tilde{y}_t(\mathbf{x}) := \psi_t(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)$ for $t \in (0, 1]$ and $\mathbf{x} \in \mathcal{X}$ where

$$\psi_t(s) = \begin{cases} \frac{64}{t} \cdot \lambda^{-1} s, & |s| \leq \frac{t}{64} \cdot \lambda, \\ \text{sign}(s), & \text{otherwise,} \end{cases} \quad s \in \mathbb{R}.$$

The absolute value of the approximation error then takes the form

$$\varepsilon_t(\mathbf{x}) = |\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \psi_t(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)| \cdot \chi_{[0, t^\circ \lambda]}(|\langle \mathbf{a}, \mathbf{x} \rangle + \tau|),$$

where we have used $t^\circ := \frac{t}{64}$ for the sake of notational convenience. We now show that for this choice of approximation, the conditions of Assumption 2 are indeed fulfilled with (8.3):

On Assumption 2(a). For $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathbb{S}^{p-1}$, we have that

$$\begin{aligned} \mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] & \leq \mathbb{E}[\chi_{[0, t^\circ \lambda]}(|\langle \mathbf{a}, \mathbf{x} \rangle + \tau|) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] \\ & \leq \mathbb{E}_{\mathbf{a}}[|\langle \mathbf{a}, \mathbf{z} \rangle| \cdot \mathbb{E}_\tau[\chi_{[0, t^\circ \lambda]}(|\langle \mathbf{a}, \mathbf{x} \rangle + \tau|)]]. \end{aligned}$$

Since $\mathbb{E}_\tau[\chi_{[0, t^\circ \lambda]}(|\langle \mathbf{a}, \mathbf{x} \rangle + \tau|)] \leq t^\circ$, it follows that

$$\mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] \leq t^\circ \cdot \mathbb{E}[|\langle \mathbf{a}, \mathbf{z} \rangle|] \leq t^\circ \cdot (\mathbb{E}[|\langle \mathbf{a}, \mathbf{z} \rangle|^2])^{1/2} = t^\circ,$$

implying the condition of Assumption 2(a).

On Assumption 2(b). Since ψ_t is $\frac{64}{t\lambda}$ -Lipschitz, the following holds for every $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$:

$$\begin{aligned} \|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} &\leq \lambda^{-1} \|\langle \mathbf{a}, \mathbf{x} - \mathbf{x}' \rangle\|_{\psi_2} + \|\psi_t(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \psi_t(\langle \mathbf{a}, \mathbf{x}' \rangle + \tau)\|_{\psi_2} \\ &\leq (\lambda^{-1} + \frac{64}{t\lambda}) \cdot \|\langle \mathbf{a}, \mathbf{x} - \mathbf{x}' \rangle\|_{\psi_2} \leq \lambda^{-1} \cdot (1 + \frac{64}{t}) \cdot L \cdot \|\mathbf{x} - \mathbf{x}'\|_2 \\ &= L \cdot (1 + 64t^{-1}) \cdot d_T(\mathbf{x}, \mathbf{x}'). \end{aligned}$$

This implies $L_t \leq L \cdot (1 + 64t^{-1})$. Furthermore, observe that $|\tilde{y}_t(\mathbf{x})| \leq 1$, so that

$$\|\xi_t(\mathbf{x})\|_{\psi_2} \leq \lambda^{-1} \|\langle \mathbf{a}, \mathbf{x} \rangle\|_{\psi_2} + \|\tilde{y}_t(\mathbf{x})\|_{\psi_2} \lesssim RL\lambda^{-1} + 1$$

for every $\mathbf{x} \in \mathcal{X}$. This shows $r \lesssim RL\lambda^{-1} + 1$.

On Assumption 2(c). We observe that the function $s \mapsto \phi_t(s) := |\text{sign}(s) - \psi_t(s)|$ is $\frac{64}{t\lambda}$ -Lipschitz. Therefore, for every $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, we obtain

$$\|\varepsilon_t(\mathbf{x}) - \varepsilon_t(\mathbf{x}')\|_{\psi_2} \leq \frac{64}{t\lambda} \cdot \|\langle \mathbf{a}, \mathbf{x} - \mathbf{x}' \rangle\|_{\psi_2} \leq 64Lt^{-1} \cdot d_T(\mathbf{x}, \mathbf{x}').$$

Hence, $\hat{L}_t \leq 64Lt^{-1}$. Finally, since $|\varepsilon_t(\mathbf{x})| \leq 1$ for every $\mathbf{x} \in \mathcal{X}$, it follows that $\hat{r} \lesssim 1$. □

8.3 Proofs for Sect. 4.3

Proof of Corollary 4 We follow the proof template from the beginning of Sect. 8:

Step 1a. The model setup of Corollary 4 fits into Assumption 1 as follows:

- The measurement vector $\mathbf{a} \in \mathbb{R}^p$ is centered, isotropic, and sub-Gaussian with $\|\mathbf{a}\|_{\psi_2} \leq L$. The signal set \mathcal{X} is a bounded subset of \mathbb{R}^p . The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(\mathbf{a}, \mathbf{x}) := q_\delta(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)$, where τ is uniformly distributed on $[-\delta, \delta]$ and independent of \mathbf{a} . In particular, F is a random function. Moreover, the observation vector of \mathbf{x} is given by $\tilde{\mathbf{y}}(\mathbf{x}) = q_\delta(\mathbf{A}\mathbf{x} + \boldsymbol{\tau})$, where $\mathbf{A} \in \mathbb{R}^{m \times p}$ denotes the sub-Gaussian random matrix with row vectors $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ and $\boldsymbol{\tau} := (\tau_1, \dots, \tau_m)$.
- The target function $T: \mathcal{X} \rightarrow K$ is the canonical embedding into K , i.e., $T\mathbf{x} := \mathbf{x}$. In particular, we have that $d_T(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_2$.

Step 1b. The target mismatch $\rho(\mathbf{x})$ vanishes for every $\mathbf{x} \in \mathcal{X}$.

Step 1c. If $t \leq 128\delta$, there exists an approximation $\tilde{y}_t(\mathbf{x})$ of the observation variable $\tilde{y}(\mathbf{x})$ such that the conditions of Assumption 2 are fulfilled with

$$L_t \lesssim L\delta t^{-1}, \quad \hat{L}_t \lesssim L\delta t^{-1}, \quad r \lesssim \delta, \quad \hat{r} \lesssim \delta. \tag{8.4}$$

If $t > 128\delta$, we choose any maximal $\frac{t}{256}$ -packing \mathcal{X}_t for \mathcal{X} (i.e., a maximal subset of \mathcal{X} such that $\|\mathbf{x} - \mathbf{x}'\|_2 > \frac{t}{256}$ for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}_t$ with $\mathbf{x} \neq \mathbf{x}'$) and show that for

the trivial choice $\tilde{y}_t(\mathbf{x}) := \tilde{\mathbf{y}}(\mathbf{x})$, the conditions of Assumption 2 are fulfilled on \mathcal{X}_t with

$$L_t \lesssim \delta t^{-1}, \quad \hat{L}_t = 0, \quad r \lesssim \delta, \quad \hat{r} = 0. \tag{8.5}$$

Step 2. We first note that for $t \geq \delta$, the inequality $c_0 L^{-2} \sqrt{\delta t} / \sqrt{\max\{1, \log(\delta e/t)\}} \leq \frac{t}{40}$ holds if $c_0 \leq \frac{1}{40}$. Furthermore, the condition $\frac{1}{m} \|\mathbf{v}\|_0 \leq t \delta^{-1}$ is trivially fulfilled. Consequently, it suffices to prove the statement of Corollary 4 for all those input vectors $\mathbf{y} \in \delta \mathbb{Z}^m$ that satisfy the condition (a) in (4.7), and all those $\mathbf{y} \in \delta \mathbb{Z}^m$ that satisfy condition (b) in (4.7) in the case $t \leq \delta$. We now distinguish between the two cases $t \leq 128\delta$ and $t > 128\delta$:

The case $t \leq 128\delta$. From (8.4), it follows that the first and third branch of the condition (2.3) is implied by (4.6) for a sufficiently large universal constant $C' > 0$. For the choice $\Delta := L^{-1} \sqrt{\log L}$, $m_0 := 0$, and $u := u_0$, we observe that the second branch of (2.3) follows from the first one. Therefore, the claim of Corollary 4 follows from Theorem 2 for all those input vectors $\mathbf{y} \in \delta \mathbb{Z}^m$ that satisfy the condition (a) in (4.7).

Next, let us assume that $t \leq \delta$ and consider the choice

$$\Delta^2 := \frac{\delta}{4CtL^4 \log(\delta e/t)}, \quad m_0 := \lfloor mt\delta^{-1} \rfloor, \quad u_0 := \sqrt{2mt\delta^{-1} \log(\delta e/t)},$$

where $C > 0$ denotes the universal constant from (2.3). This implies $u_0 \geq \sqrt{m_0 \log(em/m_0)}$ as well as $u_0^2 \geq u^2$ if the universal constant C' in (4.6) is large enough. Furthermore, the second branch of (2.3) is satisfied: We have that $CL^4 \Delta^2 u_0^2 = \frac{m}{2}$, and assuming that (4.6) holds with $C' \geq \frac{1}{2}$, it follows that

$$CL^4 \Delta^2 \cdot w_t^2(K - T\mathcal{X}) \leq \frac{1}{4} \delta t^{-1} \cdot w_t^2(K - \mathcal{X}) \leq \frac{1}{4} L^2 \delta^2 t^{-2} \cdot w_t^2(K - \mathcal{X}) \leq \frac{m}{2}.$$

It remains to show that every $\mathbf{y} \in \delta \mathbb{Z}^m$ given by (4.5) such that

$$\frac{1}{\sqrt{m}} \|\mathbf{v}\|_2 \leq \frac{c_0 \sqrt{\delta t}}{L^2 \sqrt{\log(\delta e/t)}} \quad \text{and} \quad \frac{1}{m} \|\mathbf{v}\|_0 \leq \frac{t}{\delta}$$

also satisfies the condition (2.4). Indeed, since $\mathbf{v} = \mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}) \in \delta \mathbb{Z}^m$ and $\|\mathbf{v}\|_0 \leq mt\delta^{-1}$, it follows that $|\{i \in [m] : y_i \neq \tilde{y}_i(\mathbf{x})\}| \leq \lfloor mt\delta^{-1} \rfloor$. Consequently, $\sigma_{m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}))_2 = 0$ and

$$\frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_{[m_0]} = \frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_2 \leq \frac{c_0 \sqrt{\delta t}}{L^2 \sqrt{\log(\delta e/t)}} \leq \Delta t,$$

where the last inequality holds for $c_0 > 0$ sufficiently small. Theorem 2 now implies the claim of Corollary 4 for all those input vectors $\mathbf{y} \in \delta \mathbb{Z}^m$ that satisfy the condition (b) in (4.7).

The case $t > 128\delta$. Setting $m_0 := 0$, $\Delta := L^{-1} \sqrt{\log L}$, and $u := u_0$, the second branch of (2.3) is implied by the first one. Due to (8.5), Theorem 2 applied to the

signal set \mathcal{X}_t now implies that there exist universal constants $c_1, C_1 > 0$ such that if

$$m \geq C_1 \cdot L^2 \cdot \left((\log L + t^{-2}\delta^2) \cdot (w_t^2(K - \mathcal{X}_t) + u^2) + t^{-4}\delta^2 \cdot w^2(\mathcal{X}_t) \right), \tag{8.6}$$

the following event \mathcal{A}_1 occurs with probability at least $1 - \exp(-c_1 u^2)$: For every $\mathbf{x}' \in \mathcal{X}_t$ and $\mathbf{y} \in \delta\mathbb{Z}^m$ such that $\frac{1}{\sqrt{m}} \|\mathbf{y} - q_\delta(\mathbf{A}\mathbf{x}' + \boldsymbol{\tau})\|_2 \leq \frac{t}{20}$, every minimizer $\hat{\mathbf{z}}$ of $(\mathbf{P}_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \mathbf{x}'\|_2 \leq t$.

Moreover, according to Theorem 6, there exists a universal constant $C_2 > 0$ such that if

$$m \geq C_2 \cdot L^2 \log L \cdot (w_t^2(\mathcal{X} - \mathcal{X}) + u^2), \tag{8.7}$$

the following event \mathcal{A}_2 occurs with probability at least $1 - 3 \exp(-u^2)$:

$$\sup_{\substack{\mathbf{x} \in \mathcal{X}, \mathbf{x}' \in \mathcal{X}_t \\ \|\mathbf{x} - \mathbf{x}'\|_2 \leq \frac{t}{256}}} \left\| \frac{1}{\sqrt{m}} \mathbf{A}(\mathbf{x} - \mathbf{x}') \right\|_2 \leq \frac{t}{128}.$$

We claim that Corollary 4 holds on the intersection of the events \mathcal{A}_1 and \mathcal{A}_2 . To this end, let $\mathbf{x} \in \mathcal{X}$ be arbitrary and assume that $\mathbf{y} \in \delta\mathbb{Z}^m$ satisfies $\frac{1}{\sqrt{m}} \|\mathbf{y} - q_\delta(\mathbf{A}\mathbf{x} + \boldsymbol{\tau})\|_2 \leq \frac{t}{40}$. Since \mathcal{X}_t is a maximal $\frac{t}{256}$ -packing for \mathcal{X} , there exists $\mathbf{x}' \in \mathcal{X}_t$ with $\|\mathbf{x} - \mathbf{x}'\|_2 \leq \frac{t}{256}$ (otherwise \mathcal{X}_t would not be maximal). On the event \mathcal{A}_2 and by the triangle inequality, we obtain

$$\begin{aligned} \frac{1}{\sqrt{m}} \|\mathbf{y} - q_\delta(\mathbf{A}\mathbf{x}' + \boldsymbol{\tau})\|_2 &\leq \frac{1}{\sqrt{m}} \|\mathbf{y} - q_\delta(\mathbf{A}\mathbf{x} + \boldsymbol{\tau})\|_2 \\ &\quad + \frac{1}{\sqrt{m}} \|q_\delta(\mathbf{A}\mathbf{x} + \boldsymbol{\tau}) - q_\delta(\mathbf{A}\mathbf{x}' + \boldsymbol{\tau})\|_2 \\ &\leq \frac{t}{40} + \frac{1}{\sqrt{m}} \|(\mathbf{A}\mathbf{x} + \boldsymbol{\tau}) - (\mathbf{A}\mathbf{x}' + \boldsymbol{\tau})\|_2 \\ &\quad + \frac{1}{\sqrt{m}} \|q_\delta(\mathbf{A}\mathbf{x} + \boldsymbol{\tau}) - (\mathbf{A}\mathbf{x} + \boldsymbol{\tau})\|_2 \\ &\quad + \frac{1}{\sqrt{m}} \|q_\delta(\mathbf{A}\mathbf{x}' + \boldsymbol{\tau}) - (\mathbf{A}\mathbf{x}' + \boldsymbol{\tau})\|_2 \\ &\leq \frac{t}{40} + \frac{t}{128} + 2\delta \leq \frac{t}{20}. \end{aligned}$$

On the event \mathcal{A}_1 , every minimizer $\hat{\mathbf{z}}$ of $(\mathbf{P}_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \mathbf{x}'\|_2 \leq t$, and therefore

$$\|\hat{\mathbf{z}} - \mathbf{x}\|_2 \leq \|\hat{\mathbf{z}} - \mathbf{x}'\|_2 + \|\mathbf{x}' - \mathbf{x}\|_2 \leq t + \frac{t}{256} \leq 2t.$$

Finally, for $C' > 0$ sufficiently large, we conclude that the condition (4.6) implies both (8.6) and (8.7). Hence, the events \mathcal{A}_1 and \mathcal{A}_2 occur jointly with probability at least $1 - \exp(-cu^2)$, provided that $c > 0$ is chosen small enough.

Step 3. Let $\mathbf{x} \in \mathcal{X}$. Since \mathbf{a} is isotropic, we have that

$$\rho(\mathbf{x}) = \|\mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] - \mathbf{x}\|_2 \leq \|\mathbb{E}[(q_\delta(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \mathbf{x} \rangle)\mathbf{a}]\|_2,$$

so that it suffices to show that

$$\mathbb{E}[(q_\delta(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \mathbf{x} \rangle)\langle \mathbf{a}, \mathbf{z} \rangle] = 0 \quad \text{for all } \mathbf{z} \in \mathbb{S}^{p-1}.$$

It is straightforward to check that

$$\mathbb{E}[q_\delta(s + \tau)] = s \quad \text{for all } s \in \mathbb{R}. \tag{8.8}$$

In other words, integrating over the dithering variable τ allows us to “eliminate” the quantizer. For the sake of completeness, we show the identity (8.8) for $s = k \cdot (2\delta) + s'$ where $k \in \mathbb{Z}$ and $s' \in [0, \delta]$; the case $s = k \cdot (2\delta) + s'$ where $k \in \mathbb{Z}$ and $s' \in (\delta, 2\delta)$ can be treated analogously. First, observe that

$$q_\delta(s + \tau) = (2\lceil \frac{s+\tau}{2\delta} \rceil - 1)\delta = (2\lceil k + \frac{s'+\tau}{2\delta} \rceil - 1)\delta.$$

Since $\frac{s'+\tau}{2\delta} \in [-\frac{1}{2}, 1]$, it follows that $q_\delta(s + \tau) \in \{(2k - 1)\delta, (2k + 1)\delta\}$ and more precisely,

$$\begin{aligned} \tau \in [-\delta - s', -s'] &\Rightarrow q_\delta(s + \tau) = (2k - 1)\delta, \\ \tau \in (-s', 2\delta - s'] &\Rightarrow q_\delta(s + \tau) = (2k + 1)\delta. \end{aligned}$$

Therefore, we obtain

$$\begin{aligned} \mathbb{E}[q_\delta(s + \tau)] &= \mathbb{P}(\tau \in [-\delta - s', -s']) \cdot (2k - 1)\delta \\ &\quad + \mathbb{P}(\tau \in (-s', 2\delta - s']) \cdot (2k + 1)\delta \\ &= \mathbb{P}(\tau \in [-\delta, -s']) \cdot (2k - 1)\delta + \mathbb{P}(\tau \in (-s', \delta]) \cdot (2k + 1)\delta \\ &= \left(\frac{-s'+\delta}{2\delta}\right) \cdot (2k - 1)\delta + \left(\frac{\delta+s'}{2\delta}\right) \cdot (2k + 1)\delta \\ &= \frac{s'}{2} + \frac{1}{2}(2k - 1)\delta + \frac{s'}{2} + \frac{1}{2}(2k + 1)\delta = 2k\delta + s' = s. \end{aligned}$$

Since τ and \mathbf{a} are independent, we can apply (8.8) as follows:

$$\begin{aligned} \mathbb{E}[(q_\delta(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \mathbf{x} \rangle)\langle \mathbf{a}, \mathbf{z} \rangle] &= \mathbb{E}_\mathbf{a} \mathbb{E}_\tau [(q_\delta(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \mathbf{x} \rangle)\langle \mathbf{a}, \mathbf{z} \rangle] \\ &= \mathbb{E}[(\langle \mathbf{a}, \mathbf{x} \rangle - \langle \mathbf{a}, \mathbf{x} \rangle)\langle \mathbf{a}, \mathbf{z} \rangle] = 0. \end{aligned}$$

Step 4. Before distinguishing between the cases $t \leq 128\delta$ and $t > 128\delta$ according to Step 1b, let us analyze the action of the quantizer q_δ in more detail. To this end, we partition the real axis \mathbb{R} into half-open intervals \mathcal{I}_k of length 2δ given by

$\mathcal{I}_k := (e_{k-1}, e_k]$ for $k \in \mathbb{Z}$ with $e_k := k \cdot 2\delta$. Then, q_δ maps every point in the interval \mathcal{I}_k to its center point:

$$q_\delta(s) = \frac{e_{k-1} + e_k}{2}, \quad s \in \mathcal{I}_k.$$

In particular, $q_\delta(s)$ is discontinuous exactly at the points $\{e_k : k \in \mathbb{Z}\}$. The basic idea is now to approximate the quantizer q_δ by a Lipschitz continuous function by “cutting out” intervals of a certain radius $t^\circ > 0$ around each point e_k and inserting a straight line that connects both quantization values. Assuming that $t^\circ \leq \delta$, the resulting function takes the form

$$\psi_{t^\circ}(s) := \begin{cases} \frac{\delta}{t^\circ} \cdot s + e_k(1 - \frac{\delta}{t^\circ}), & \text{if } \exists k \in \mathbb{Z} : s \in \mathcal{E}_{k,t^\circ}, \\ q_\delta(s), & \text{otherwise,} \end{cases} \quad s \in \mathbb{R},$$

where $\mathcal{E}_{k,t^\circ} := [e_k - t^\circ, e_k + t^\circ]$; note that if $t^\circ = \delta$, then ψ_{t° is just the identity function. We also define

$$\phi_{t^\circ}(s) := |q_\delta(s) - \psi_{t^\circ}(s)|, \quad s \in \mathbb{R}.$$

Let us now show that ψ_{t° and ϕ_{t° are both $\frac{\delta}{t^\circ}$ -Lipschitz.

Since q_δ is locally constant on $\mathbb{R} \setminus \{e_k : k \in \mathbb{Z}\}$ and $s \mapsto \frac{\delta}{t^\circ} \cdot s + e_k(1 - \frac{\delta}{t^\circ})$ is clearly $\frac{\delta}{t^\circ}$ -Lipschitz, it is sufficient to show that ψ_{t° is continuous, i.e., $\frac{\delta}{t^\circ} \cdot s + e_k(1 - \frac{\delta}{t^\circ}) = q_\delta(s)$ for $s \in \mathbb{R}$ with $|s - e_k| = t^\circ$ for some $k \in \mathbb{Z}$. If $s = e_k - t^\circ$, then $\frac{\delta}{t^\circ} \cdot s + e_k(1 - \frac{\delta}{t^\circ}) = e_k - \delta$, and we also have that

$$q_\delta(s) = (2 \lceil \frac{k \cdot (2\delta) - t^\circ}{2\delta} \rceil - 1)\delta = (2 \lceil k - \frac{t^\circ}{2\delta} \rceil - 1)\delta.$$

Since $t^\circ \in (0, \delta]$, it follows that $\lceil k - \frac{t^\circ}{2\delta} \rceil = k$, and therefore, $q_\delta(s) = (2k - 1)\delta = e_k - \delta$. The case $s = e_k + t^\circ$ works analogously. Thus, ψ_{t° is indeed $\frac{\delta}{t^\circ}$ -Lipschitz. For ϕ_{t° , it is clearly sufficient to show $\frac{\delta}{t^\circ}$ -Lipschitz continuity on \mathcal{E}_{k,t° for every $k \in \mathbb{Z}$. In the case $s \in [e_k - t^\circ, e_k]$, we observe that $q_\delta(s) = e_k - \delta$ and $\frac{\delta}{t^\circ} \cdot s + e_k(1 - \frac{\delta}{t^\circ}) \geq e_k - \delta$, implying that

$$\phi_{t^\circ}(s) = \frac{\delta}{t^\circ} \cdot s + e_k(1 - \frac{\delta}{t^\circ}) - (e_k - \delta) = \delta - \frac{\delta}{t^\circ}(e_k - s).$$

Thus, ϕ_{t° is $\frac{\delta}{t^\circ}$ -Lipschitz continuous on $[e_k - t^\circ, e_k]$. On the other hand, if $s \in (e_k, e_k + t^\circ]$, we have that $q_\delta(s) = e_k + \delta$ and $e_k + \delta \geq \frac{\delta}{t^\circ} \cdot s + e_k(1 - \frac{\delta}{t^\circ})$, implying that

$$\phi_{t^\circ}(s) = e_k + \delta - (\frac{\delta}{t^\circ} \cdot s + e_k(1 - \frac{\delta}{t^\circ})) = \delta - \frac{\delta}{t^\circ}(s - e_k).$$

Thus, ϕ_{t° is $\frac{\delta}{t^\circ}$ -Lipschitz continuous on $(e_k, e_k + t^\circ]$. Moreover, $\lim_{s \rightarrow e_k} \phi_{t^\circ}(s) = \delta = \phi_{t^\circ}(e_k)$, which shows that ϕ_{t° is indeed $\frac{\delta}{t^\circ}$ -Lipschitz continuous on \mathcal{E}_{k,t° . Note

that the above calculations also allow us to write

$$\phi_{t^\circ}(s) = \begin{cases} \delta - \frac{\delta}{t^\circ} |s - e_k|, & \text{if } \exists k \in \mathbb{Z} : s \in [e_k - t^\circ, e_k + t^\circ], \\ 0, & \text{otherwise,} \end{cases} \quad s \in \mathbb{R}.$$

In particular, we have that $0 \leq \phi_{t^\circ}(s) \leq \delta \cdot \chi_{\mathcal{E}_{t^\circ}}(s)$, where $\mathcal{E}_{t^\circ} := \cup_{k \in \mathbb{Z}} \mathcal{E}_{k, t^\circ}$. We now distinguish between the two cases $t \leq 128\delta$ and $t > 128\delta$:

The case $t \leq 128\delta$. Using the above notation, we choose $t^\circ := \frac{t}{128}$ and approximate $\tilde{y}(\mathbf{x}) = q_\delta(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)$ by $\tilde{y}_t(\mathbf{x}) := \psi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)$. The absolute value of the approximation error is then given by $\varepsilon_t(\mathbf{x}) = \phi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)$. We now show that for this choice of approximation, the conditions of Assumption 2 are indeed fulfilled with (8.4):

On Assumption 2(a). We first observe that the indicator function $\chi_{\mathcal{E}_{t^\circ}}$ is 2δ -periodic on \mathbb{R} . Since the random variable $s + \tau$ is uniformly distributed on $[s - \delta, s + \delta]$ for every $s \in \mathbb{R}$, this implies

$$\mathbb{E}_\tau[\chi_{\mathcal{E}_{t^\circ}}(s + \tau)] = \mathbb{E}_\tau[\chi_{\mathcal{E}_{t^\circ}}(\tau)] = \mathbb{E}_\tau[\chi_{\mathcal{E}_{0, t^\circ}}(\tau)] = \mathbb{E}_\tau[\chi_{[-t^\circ, t^\circ]}(\tau)] = t^\circ \delta^{-1} \tag{8.9}$$

Now, let $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathbb{S}^{p-1}$. Using the independence of \mathbf{a} and τ in conjunction with the inequality (8.9), we obtain

$$\begin{aligned} \mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] &= \mathbb{E}[\phi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] \\ &\leq \delta \cdot \mathbb{E}[\chi_{\mathcal{E}_{t^\circ}}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] \\ &\leq \delta \cdot \mathbb{E}_\mathbf{a}[|\langle \mathbf{a}, \mathbf{z} \rangle| \cdot \mathbb{E}_\tau[\chi_{\mathcal{E}_{t^\circ}}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)]] \\ &\leq t^\circ \cdot \mathbb{E}[|\langle \mathbf{a}, \mathbf{z} \rangle|] \leq \frac{t}{64} \cdot \mathbb{E}[|\langle \mathbf{a}, \mathbf{z} \rangle|]. \end{aligned}$$

From Jensen’s inequality and the isotropy of \mathbf{a} , it follows that

$$\mathbb{E}[|\langle \mathbf{a}, \mathbf{z} \rangle|] \leq (\mathbb{E}[|\langle \mathbf{a}, \mathbf{z} \rangle|^2])^{1/2} = \|\mathbf{z}\|_2 = 1,$$

which shows that Assumption 2(a) is satisfied.

On Assumption 2(b). Since ψ_{t° is $\frac{\delta}{t^\circ}$ -Lipschitz, the following holds for every $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$:

$$\begin{aligned} \|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} &\leq \|\langle \mathbf{a}, \mathbf{x} - \mathbf{x}' \rangle\|_{\psi_2} + \|\psi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \psi_{t^\circ}(\langle \mathbf{a}, \mathbf{x}' \rangle + \tau)\|_{\psi_2} \\ &\leq (1 + \frac{\delta}{t^\circ}) \cdot \|\langle \mathbf{a}, \mathbf{x} - \mathbf{x}' \rangle\|_{\psi_2} \leq (1 + \frac{\delta}{t^\circ}) \cdot L \cdot \|\mathbf{x} - \mathbf{x}'\|_2 \\ &= L \cdot (1 + \frac{128\delta}{t}) \cdot d_T(\mathbf{x}, \mathbf{x}'). \end{aligned}$$

This implies $L_t \lesssim L\delta t^{-1}$. Furthermore, we clearly have that $\sup_{s \in \mathbb{R}} |s - \psi_{t^\circ}(s + \tau)| \lesssim \delta$, and therefore

$$\|\xi_t(\mathbf{x})\|_{\psi_2} = \|\langle \mathbf{a}, \mathbf{x} \rangle - \psi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)\|_{\psi_2} \lesssim \delta$$

for every $\mathbf{x} \in \mathcal{X}$. This shows $r \lesssim \delta$.

On Assumption 2(c). Since ϕ_{t° is $\frac{\delta}{t^\circ}$ -Lipschitz, the following holds for every $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$:

$$\begin{aligned} \|\varepsilon_t(\mathbf{x}) - \varepsilon_t(\mathbf{x}')\|_{\psi_2} &= \|\phi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \phi_{t^\circ}(\langle \mathbf{a}, \mathbf{x}' \rangle + \tau)\|_{\psi_2} \\ &\leq \frac{\delta}{t^\circ} \|\langle \mathbf{a}, \mathbf{x} - \mathbf{x}' \rangle\|_{\psi_2} \leq L \cdot \frac{128\delta}{t} \cdot d_T(\mathbf{x}, \mathbf{x}'), \end{aligned}$$

and therefore, $\hat{L}_t \lesssim L\delta t^{-1}$. Finally, since $0 \leq \phi_{t^\circ}(s) \leq \delta$ for every $s \in \mathbb{R}$, it follows that $\hat{r} \lesssim \delta$.

The case $t > 128\delta$. Let $\mathcal{X}_t \subset \mathcal{X}$ be any maximal $\frac{t}{256}$ -packing for \mathcal{X} . We show that for the choice $\tilde{y}_t(\mathbf{x}) := \tilde{y}(\mathbf{x})$, the conditions of Assumption 2 are fulfilled on \mathcal{X}_t with (8.5). Since $\varepsilon_t(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathcal{X}_t$, Assumption 2(a) and (c) are trivially fulfilled with $\hat{L}_t = \hat{r} = 0$. Furthermore, observing that $\xi_t(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle - q_\delta(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)$ and $\sup_{s \in \mathbb{R}} |s - q_\delta(s + \tau)| \leq 2\delta$, we conclude that $\sup_{\mathbf{x} \in \mathcal{X}_t} \|\xi_t(\mathbf{x})\|_{\psi_2} \lesssim \delta$. Finally, using that $\|\mathbf{x} - \mathbf{x}'\|_2 > \frac{t}{256}$ for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}_t$ with $\mathbf{x} \neq \mathbf{x}'$, we can bound the sub-Gaussian norm of the multiplier increments as follows:

$$\begin{aligned} \|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} &\leq \|\xi_t(\mathbf{x})\|_{\psi_2} + \|\xi_t(\mathbf{x}')\|_{\psi_2} \lesssim \delta < \frac{256\delta}{t} \|\mathbf{x} - \mathbf{x}'\|_2 \\ &= \frac{256\delta}{t} \cdot d_T(\mathbf{x}, \mathbf{x}'). \end{aligned}$$

This shows that Assumption 2(b) is satisfied on \mathcal{X}_t with $L_t \lesssim \delta t^{-1}$ and $r \lesssim \delta$. □

8.4 Proofs for Sect. 4.4

Proof of Theorem 1 We follow the proof template from the beginning of Sect. 8:

Step 1a. The model setup of Theorem 1 fits into Assumption 1 as follows:

- We have that $\mathbf{a} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ and therefore $\|\mathbf{a}\|_{\psi_2} \lesssim 1$. The signal set \mathcal{X} is an arbitrary subset of \mathbb{S}^{p-1} . The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(\mathbf{a}, \mathbf{x}) := f(\langle \mathbf{a}, \mathbf{x} \rangle)$.
- The target function $T: \mathcal{X} \rightarrow K$ corresponds to rescaling by a factor of $\mu = \mathbb{E}[f(g)g]$ with $g \sim \mathcal{N}(0, 1)$, i.e., $T\mathbf{x} := \mu\mathbf{x}$. In particular, we have that $d_T(\mathbf{x}, \mathbf{x}') = \mu\|\mathbf{x} - \mathbf{x}'\|_2$.

Step 1b. The target mismatch $\rho(\mathbf{x})$ vanishes for every $\mathbf{x} \in \mathcal{X}$.

Step 1c. Let $g \sim \mathcal{N}(0, 1)$. For the trivial choice $\tilde{y}_t(\mathbf{x}) := \tilde{y}(\mathbf{x})$, the conditions of Assumption 2 are fulfilled with

$$L_t \lesssim 1 + \gamma\mu^{-1}, \quad \hat{L}_t = 0, \quad r = \|f(g) - \mu g\|_{\psi_2}, \quad \hat{r} = 0.$$

Step 2. Setting $m_0 := 0$, $\Delta := 1$, and $u_0 := u$, the claim of Theorem 1 follows directly from Theorem 2.

Step 3. Let $\mathbf{x} \in \mathcal{X} \subset \mathbb{S}^{p-1}$ and consider the orthogonal decomposition of the standard Gaussian random vector \mathbf{a} along \mathbf{x} :

$$\mathbf{a} = \langle \mathbf{a}, \mathbf{x} \rangle \mathbf{x} + P_{\mathbf{x}^\perp}(\mathbf{a}),$$

where $P_{\mathbf{x}^\perp} := P_{\{\mathbf{x}\}^\perp}$. Since $P_{\mathbf{x}^\perp}(\mathbf{a})$ is centered and $\langle \mathbf{a}, \mathbf{x} \rangle \sim \mathcal{N}(0, 1)$ is independent of $P_{\mathbf{x}^\perp}(\mathbf{a})$, we have that

$$\begin{aligned} \mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] &= \mathbb{E}[f(\langle \mathbf{a}, \mathbf{x} \rangle)(\langle \mathbf{a}, \mathbf{x} \rangle \mathbf{x} + P_{\mathbf{x}^\perp}(\mathbf{a}))] \\ &= \mu \mathbf{x} + \mathbb{E}[f(\langle \mathbf{a}, \mathbf{x} \rangle)] \cdot \mathbb{E}[P_{\mathbf{x}^\perp}(\mathbf{a})] = \mu \mathbf{x} = T\mathbf{x}, \end{aligned}$$

which implies that $\rho(\mathbf{x}) = 0$.

Step 4. We simply set $\tilde{\mathbf{y}}_t(\mathbf{x}) := \tilde{\mathbf{y}}(\mathbf{x})$. Then, $\varepsilon_t(\mathbf{x}) = 0$, implying that Assumption 2(a) and (c) are trivially fulfilled with $\hat{L}_t = \hat{r} = 0$. Furthermore, the following holds for every $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$:

$$\begin{aligned} \|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} &\leq \mu \|\langle \mathbf{a}, \mathbf{x} - \mathbf{x}' \rangle\|_{\psi_2} + \|f(\langle \mathbf{a}, \mathbf{x} \rangle) - f(\langle \mathbf{a}, \mathbf{x}' \rangle)\|_{\psi_2} \\ &\lesssim \mu \|\mathbf{x} - \mathbf{x}'\|_2 + \gamma \|\langle \mathbf{a}, \mathbf{x} - \mathbf{x}' \rangle\|_{\psi_2} \\ &\lesssim (\mu + \gamma) \cdot \|\mathbf{x} - \mathbf{x}'\|_2 = (1 + \gamma\mu^{-1}) \cdot d_T(\mathbf{x}, \mathbf{x}'). \end{aligned}$$

This implies $L_t \lesssim 1 + \gamma\mu^{-1}$. Since $\langle \mathbf{a}, \mathbf{x} \rangle \sim \mathcal{N}(0, 1)$ for every $\mathbf{x} \in \mathcal{X} \subset \mathbb{S}^{p-1}$, we can also conclude that

$$\|\xi_t(\mathbf{x})\|_{\psi_2} = \|\mu \langle \mathbf{a}, \mathbf{x} \rangle - f(\langle \mathbf{a}, \mathbf{x} \rangle)\|_{\psi_2} = \|f(g) - \mu g\|_{\psi_2}, \quad g \sim \mathcal{N}(0, 1),$$

which shows that $r := \|f(g) - \mu g\|_{\psi_2}$ is a valid choice. Hence, Assumption 2(b) is satisfied as well. □

Proof of Corollary 5 We follow the proof template from the beginning of Sect. 8:

Step 1a. The model setup of Corollary 5 fits into Assumption 1 as follows:

- We have that $\mathbf{a} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ and therefore $\|\mathbf{a}\|_{\psi_2} \lesssim 1$. The signal set \mathcal{X} is an arbitrary subset of \mathbb{S}^{p-1} . The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(\mathbf{a}, \mathbf{x}) := m_\lambda(\langle \mathbf{a}, \mathbf{x} \rangle)$.
- The target function $T: \mathcal{X} \rightarrow K$ corresponds to rescaling by a factor of $\mu_\lambda = \mathbb{E}[m_\lambda(g)g]$ with $g \sim \mathcal{N}(0, 1)$, i.e., $T\mathbf{x} := \mu_\lambda \mathbf{x}$. In particular, we have that $d_T(\mathbf{x}, \mathbf{x}') = \mu_\lambda \|\mathbf{x} - \mathbf{x}'\|_2$.

Step 1b. The target mismatch $\rho(\mathbf{x})$ vanishes for every $\mathbf{x} \in \mathcal{X}$.

Step 1c. There exists an approximation $\tilde{\mathbf{y}}_t(\mathbf{x})$ of the observation variable $\tilde{\mathbf{y}}(\mathbf{x})$ such that the conditions of Assumption 2 are fulfilled with

$$L_t \lesssim \frac{\lambda}{t}, \quad \hat{L}_t \lesssim \frac{\lambda}{t}, \quad r \lesssim 1, \quad \hat{r} \lesssim 1. \tag{8.10}$$

Step 2. We begin by showing that if $\lambda \geq C$ for a sufficiently large absolute constant $C > 0$, then $\mu_\lambda \geq \frac{1}{2}$. Indeed, for $g \sim N(0, 1)$, we have that

$$\begin{aligned} \mu_\lambda &= \mathbb{E}[m_\lambda(g)g] = \mathbb{E}[\chi_{[-\lambda,\lambda]}(g)g^2] + \mathbb{E}[\chi_{\mathbb{R}\setminus[-\lambda,\lambda]}(g)m_\lambda(g)g] \\ &\geq \mathbb{E}[\chi_{[-\lambda,\lambda]}(g)g^2] - \mathbb{E}[\chi_{\mathbb{R}\setminus[-\lambda,\lambda]}(g)g^2] \\ &= \mathbb{E}[g^2] - 2\mathbb{E}[\chi_{\mathbb{R}\setminus[-\lambda,\lambda]}(g)g^2] \\ &= 1 - 2\mathbb{E}[\chi_{\mathbb{R}\setminus[-\lambda,\lambda]}(g)g^2]. \end{aligned}$$

The claim now follows by observing that $\mathbb{E}[g^2] < \infty$, which implies $\lim_{\lambda \rightarrow \infty} \mathbb{E}[\chi_{\mathbb{R}\setminus[-\lambda,\lambda]}(g)g^2] = 0$. Moreover,

$$\begin{aligned} \mu_\lambda &= \mathbb{E}[\chi_{[-\lambda,\lambda]}(g)g^2] + \mathbb{E}[\chi_{\mathbb{R}\setminus[-\lambda,\lambda]}(g)m_\lambda(g)g] \\ &\leq \mathbb{E}[\chi_{[-\lambda,\lambda]}(g)g^2] + \mathbb{E}[\chi_{\mathbb{R}\setminus[-\lambda,\lambda]}(g)\lambda|g|] \leq \mathbb{E}[g^2] = 1. \end{aligned}$$

Setting $m_0 := 0$, $\Delta := 1$, and $u_0 := u$, the claim of Corollary 5 follows directly from Theorem 2.

Step 3. Let $\mathbf{x} \in \mathcal{X} \subset \mathbb{S}^{p-1}$ and consider the orthogonal decomposition of the standard Gaussian random vector \mathbf{a} along \mathbf{x} :

$$\mathbf{a} = \langle \mathbf{a}, \mathbf{x} \rangle \mathbf{x} + P_{\mathbf{x}^\perp}(\mathbf{a}),$$

where $P_{\mathbf{x}^\perp} := P_{\{\mathbf{x}\}^\perp}$. Since $P_{\mathbf{x}^\perp}(\mathbf{a})$ is centered and $\langle \mathbf{a}, \mathbf{x} \rangle \sim N(0, 1)$ is independent of $P_{\mathbf{x}^\perp}(\mathbf{a})$, we have that

$$\begin{aligned} \mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] &= \mathbb{E}[m_\lambda(\langle \mathbf{a}, \mathbf{x} \rangle)(\langle \mathbf{a}, \mathbf{x} \rangle \mathbf{x} + P_{\mathbf{x}^\perp}(\mathbf{a}))] \\ &= \mu_\lambda \mathbf{x} + \mathbb{E}[m_\lambda(\langle \mathbf{a}, \mathbf{x} \rangle)] \cdot \mathbb{E}[P_{\mathbf{x}^\perp}(\mathbf{a})] = \mu_\lambda \mathbf{x} = T\mathbf{x}, \end{aligned}$$

which implies that $\rho(\mathbf{x}) = 0$.

Step 4.

Let $t \in (0, \lambda)$. We define the half-open intervals $\mathcal{I}_{k,t} := [k\lambda - (\lambda - t), k\lambda + (\lambda - t))$ for $k \in \mathbb{Z}$ even and $\mathcal{I}_{k,t} := [k\lambda - t, k\lambda + t)$ for $k \in \mathbb{Z}$ odd. Moreover, we set

$$\mathcal{I}_{\text{even},t} := \bigcup_{\substack{k \in \mathbb{Z} \\ k \text{ even}}} \mathcal{I}_{k,t}, \quad \mathcal{I}_{\text{odd},t} := \bigcup_{\substack{k \in \mathbb{Z} \\ k \text{ odd}}} \mathcal{I}_{k,t}$$

and observe that the intervals $\mathcal{I}_{k,t}$, $k \in \mathbb{Z}$, form a partition of \mathbb{R} . Finally, we introduce the functions

$$g_t(s) := -\left(\frac{\lambda-t}{t}\right) \cdot s, \quad s \in \mathbb{R},$$

and

$$\psi_t(s) := \begin{cases} m_\lambda(s), & s \in \mathcal{I}_{k,t}, \quad k \text{ even}, \\ g_t(s - k\lambda), & s \in \mathcal{I}_{k,t}, \quad k \text{ odd}, \end{cases}$$

and

$$\phi_t(s) := |m_\lambda(s) - \psi_t(s)|, \quad s \in \mathbb{R}.$$

Let us show that both ψ_t and ϕ_t are $\frac{\lambda}{t}$ -Lipschitz continuous. We begin by verifying that ψ_t is continuous. Clearly, ψ_t is continuous on $\mathcal{I}_{k,t}$ for every odd integer k . Let k be an even integer. Then, $s \in \mathcal{I}_{k,t} = [k\lambda - (\lambda - t), k\lambda + (\lambda - t))$, which implies $\lfloor \frac{s+\lambda}{2\lambda} \rfloor = \frac{k}{2}$. Therefore, $m_\lambda(s) = s - k\lambda$ if $s \in \mathcal{I}_{k,t}$ for k an even integer. This shows that ψ_t is continuous on every interval $\mathcal{I}_{k,t}$. Let us now write $\mathcal{I}_{k,t} = [u_k, v_k)$ for $k \in \mathbb{Z}$. To show that ψ_t is continuous, it remains to check that $\lim_{s \uparrow v_k} \psi_t(s) = \psi_t(u_{k+1})$ for every $k \in \mathbb{Z}$. For k even, we have that

$$\lim_{s \uparrow v_k} \psi_t(s) = \lim_{s \uparrow v_k} m_\lambda(s) = \lim_{s \uparrow v_k} s - k\lambda = v_k - k\lambda = \lambda - t.$$

Furthermore, it holds that $u_{k+1} = (k + 1)\lambda - t \in \mathcal{I}_{k+1,t}$, which implies

$$\psi_t(u_{k+1}) = g_t(u_{k+1} - (k + 1)\lambda) = g_t(-t) = \lambda - t.$$

On the other hand, for k odd, we have that

$$\lim_{s \uparrow v_k} \psi_t(s) = \lim_{s \uparrow v_k} g_t(s - k\lambda) = \lim_{s \uparrow v_k} -\left(\frac{\lambda-t}{t}\right)(s - k\lambda) = -\left(\frac{\lambda-t}{t}\right)(v_k - k\lambda) = -(\lambda - t).$$

Furthermore, it holds that $u_{k+1} = k\lambda + t \in \mathcal{I}_{k+1,t}$, which implies

$$\psi_t(u_{k+1}) = m_\lambda(u_{k+1}) = u_{k+1} - (k + 1)\lambda = -(\lambda - t).$$

This shows that ψ_t is continuous. Since ψ_t is continuous and piecewise linear, it is $\max\{1, \frac{\lambda-t}{t}\}$ -Lipschitz. Therefore, ψ_t is also $\frac{\lambda}{t}$ -Lipschitz.

Next, we show that ϕ_t is $\frac{\lambda}{t}$ -Lipschitz as well. First note that

$$\phi_t(s) = |m_\lambda(s) - \psi_t(s)| \cdot \chi_{\mathcal{I}_{\text{odd},t}}(s), \quad s \in \mathbb{R}.$$

Let $s \in \mathcal{I}_{k,t} = [k\lambda - t, k\lambda + t)$ for k an odd integer. Let us show that

$$|m_\lambda(s) - \psi_t(s)| = \lambda - \frac{\lambda}{t}|k\lambda - s|.$$

To this end, we distinguish the cases $s \in [k\lambda - t, k\lambda)$ and $s \in [k\lambda, k\lambda + t)$. First, let $s \in [k\lambda - t, k\lambda)$. It follows $\lfloor \frac{s+\lambda}{2\lambda} \rfloor = \frac{k-1}{2}$, which implies $m_\lambda(s) = s - k\lambda + \lambda$. Furthermore, $g_t(s - k\lambda) = s - k\lambda + \frac{\lambda}{t}(k\lambda - s)$. It follows that $m_\lambda(s) \geq g_t(s - k\lambda)$ and

$$|m_\lambda(s) - \psi_t(s)| = m_\lambda(s) - g_t(s - k\lambda) = \lambda - \frac{\lambda}{t}|k\lambda - s|.$$

The case $s \in [k\lambda, k\lambda + t)$ can be treated analogously. All together, this shows that for any $s \in \mathbb{R}$, we have that

$$\phi_t(s) = \sum_{\substack{k \in \mathbb{Z} \\ k \text{ odd}}} \chi_{\mathcal{I}_{k,t}}(s) \cdot (\lambda - \frac{\lambda}{t} |k\lambda - s|).$$

In particular, ϕ_t is $\frac{\lambda}{t}$ -Lipschitz.

We approximate the observation variable $\tilde{y}(\mathbf{x}) = m_\lambda(\langle \mathbf{a}, \mathbf{x} \rangle)$ by $\tilde{y}_t(\mathbf{x}) = \psi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle)$ for $t^\circ := \frac{t}{56 \cdot 64}$. Observe that $t^\circ < \frac{\lambda}{2}$. Let us show that the conditions of Assumption 2 are indeed fulfilled with (8.10):

On Assumption 2(a). Let $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathbb{S}^{p-1}$, and consider the orthogonal decomposition

$$\mathbf{z} = \langle \mathbf{z}, \mathbf{x} \rangle \mathbf{x} + P_{\mathbf{x}^\perp}(\mathbf{z}).$$

This implies

$$\mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] \leq |\langle \mathbf{z}, \mathbf{x} \rangle| \cdot \mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |\langle \mathbf{a}, \mathbf{x} \rangle|] + \mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |\langle \mathbf{a}, P_{\mathbf{x}^\perp}(\mathbf{z}) \rangle|].$$

Clearly, $|\langle \mathbf{z}, \mathbf{x} \rangle| \leq 1$ and $\varepsilon_t(\mathbf{x}) = \phi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle) \leq \lambda \cdot \chi_{\mathcal{I}_{\text{odd},t^\circ}}(\langle \mathbf{a}, \mathbf{x} \rangle)$, which implies

$$\begin{aligned} \mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] &\leq \lambda \cdot \underbrace{\mathbb{E}[\chi_{\mathcal{I}_{\text{odd},t^\circ}}(\langle \mathbf{a}, \mathbf{x} \rangle) \cdot |\langle \mathbf{a}, \mathbf{x} \rangle|]}_{=: S_1} \\ &\quad + \lambda \cdot \underbrace{\mathbb{E}[\chi_{\mathcal{I}_{\text{odd},t^\circ}}(\langle \mathbf{a}, \mathbf{x} \rangle) \cdot |\langle \mathbf{a}, P_{\mathbf{x}^\perp}(\mathbf{z}) \rangle|]}_{=: S_2}. \end{aligned}$$

Let us estimate S_1 and S_2 separately, starting with S_2 . Since $\langle \mathbf{a}, \mathbf{x} \rangle$ and $\langle \mathbf{a}, P_{\mathbf{x}^\perp}(\mathbf{z}) \rangle$ are independent, we observe that

$$\mathbb{E}[\chi_{\mathcal{I}_{\text{odd},t^\circ}}(\langle \mathbf{a}, \mathbf{x} \rangle) \cdot |\langle \mathbf{a}, P_{\mathbf{x}^\perp}(\mathbf{z}) \rangle|] = \mathbb{E}[\chi_{\mathcal{I}_{\text{odd},t^\circ}}(\langle \mathbf{a}, \mathbf{x} \rangle)] \cdot \mathbb{E}[|\langle \mathbf{a}, P_{\mathbf{x}^\perp}(\mathbf{z}) \rangle|].$$

Using Jensen’s inequality, we obtain

$$\mathbb{E}[|\langle \mathbf{a}, P_{\mathbf{x}^\perp}(\mathbf{z}) \rangle|] \leq \left(\mathbb{E}[|\langle \mathbf{a}, P_{\mathbf{x}^\perp}(\mathbf{z}) \rangle|^2] \right)^{1/2} = \|P_{\mathbf{x}^\perp}(\mathbf{z})\|_2 \leq \|\mathbf{z}\|_2 = 1.$$

Therefore, it follows that $S_2 \leq \lambda \cdot \mathbb{P}(\langle \mathbf{a}, \mathbf{x} \rangle \in \mathcal{I}_{\text{odd},t^\circ})$. Since $\langle \mathbf{a}, \mathbf{x} \rangle \sim \mathbf{N}(0, 1)$, it holds that

$$\frac{1}{2} \mathbb{P}(\langle \mathbf{a}, \mathbf{x} \rangle \in \mathcal{I}_{\text{odd},t^\circ}) = \frac{1}{\sqrt{2\pi}} \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} \int_{k\lambda - t^\circ}^{k\lambda + t^\circ} e^{-x^2/2} dx \leq \frac{2t^\circ}{\sqrt{2\pi}} \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} e^{-(k\lambda - t^\circ)^2/2}.$$

For every $k \in \mathbb{N}$,

$$2\lambda e^{-((k+2)\lambda - t^\circ)^2/2} \leq \int_{k\lambda - t^\circ}^{(k+2)\lambda - t^\circ} e^{-x^2/2} dx.$$

Therefore,

$$\begin{aligned} \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} e^{-(k\lambda - t^\circ)^2/2} &= e^{-(\lambda - t^\circ)^2/2} + \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} e^{-((k+2)\lambda - t^\circ)^2/2} \\ &\leq \frac{1}{\lambda - t^\circ} \int_0^{\lambda - t^\circ} e^{-x^2/2} dx + \frac{1}{2\lambda} \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} \int_{k\lambda - t^\circ}^{(k+2)\lambda - t^\circ} e^{-x^2/2} dx \\ &\leq \frac{1}{\lambda - t^\circ} \int_0^\infty e^{-x^2/2} dx. \end{aligned}$$

Using the identity $\frac{1}{\sqrt{2\pi}} \int_0^\infty e^{-x^2/2} dx = \frac{1}{2}$, we obtain

$$S_2 \leq \frac{2t^\circ\lambda}{\lambda - t^\circ} \leq 4t^\circ,$$

where the last inequality follows from $t^\circ < \frac{\lambda}{2}$.

Let us now estimate S_1 . Since $\langle \mathbf{a}, \mathbf{x} \rangle \sim N(0, 1)$, we have that

$$\mathbb{E}[X_{\mathcal{I}_{\text{odd}, t^\circ}}(\langle \mathbf{a}, \mathbf{x} \rangle) \cdot |\langle \mathbf{a}, \mathbf{x} \rangle|] = \frac{1}{\sqrt{2\pi}} \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} \int_{k\lambda - t^\circ}^{k\lambda + t^\circ} x e^{-x^2/2} dx.$$

It holds that

$$\begin{aligned} &\sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} \int_{k\lambda - t^\circ}^{k\lambda + t^\circ} x e^{-x^2/2} dx \\ &\leq 2t^\circ \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} (k\lambda + t^\circ) e^{-(k\lambda - t^\circ)^2/2} \\ &= 2t^\circ \left((\lambda + t^\circ) \cdot e^{-(\lambda - t^\circ)^2/2} + \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} ((k+2)\lambda + t^\circ) e^{-((k+2)\lambda - t^\circ)^2/2} dx \right) \\ &\leq 2t^\circ \left(\frac{1}{\lambda - t^\circ} \int_0^{\lambda - t^\circ} (x + \lambda + t^\circ) \cdot e^{-x^2/2} dx \right. \\ &\quad \left. + \frac{1}{2\lambda} \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} \int_{k\lambda - t^\circ}^{(k+2)\lambda - t^\circ} (x + 2\lambda + 2t^\circ) \cdot e^{-x^2/2} dx \right) \\ &\leq \frac{4t^\circ}{\lambda} \int_0^\infty (x + 3\lambda) \cdot e^{-x^2/2} dx \end{aligned}$$

$$= \frac{4t^\circ}{\lambda} \cdot (1 + 3\lambda\sqrt{\frac{\pi}{2}}).$$

Therefore, if $\lambda \leq 2$, then

$$\sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} \int_{k\lambda - t^\circ}^{k\lambda + t^\circ} x e^{-x^2/2} dx \leq \frac{4t^\circ}{\lambda} \cdot (1 + 6\sqrt{\frac{\pi}{2}}) \leq \frac{52t^\circ}{\lambda}.$$

If $\lambda > 2$, then $\lambda - t^\circ > 1$. Since the function $x \mapsto x \exp(-x^2/2)$ is monotonically decreasing for $x \in [1, \infty)$, we obtain

$$\begin{aligned} & \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} \int_{k\lambda - t^\circ}^{k\lambda + t^\circ} x e^{-x^2/2} dx \\ & \leq 2t^\circ \cdot \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} (k\lambda - t^\circ) \cdot e^{-(k\lambda - t^\circ)^2/2} \\ & = 2t^\circ \cdot \left((\lambda - t^\circ) \cdot \exp(-(\lambda - t^\circ)^2/2) + \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} ((k+2)\lambda - t^\circ) \cdot e^{-((k+2)\lambda - t^\circ)^2/2} \right). \end{aligned}$$

Since $\lambda - t^\circ > 1$, it follows $(\lambda - t^\circ) \cdot \exp(-(\lambda - t^\circ)^2/2) \leq \frac{1}{\lambda - t^\circ}$. Furthermore,

$$\sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} ((k+2)\lambda - t^\circ) \cdot e^{-((k+2)\lambda - t^\circ)^2/2} \leq \frac{1}{2\lambda} \sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} \int_{k\lambda - t^\circ}^{(k+2)\lambda - t^\circ} x e^{-x^2/2} dx \leq \frac{1}{2\lambda}.$$

Therefore, if $\lambda > 2$, then

$$\sum_{\substack{k \in \mathbb{N} \\ k \text{ odd}}} \int_{k\lambda - t^\circ}^{k\lambda + t^\circ} x e^{-x^2/2} dx \leq \frac{8t^\circ}{\lambda}.$$

This shows $S_1 \leq 52t^\circ$. Putting everything together, we obtain

$$\mathbb{E}[\varepsilon_t(\mathbf{x}) \cdot |\langle \mathbf{a}, \mathbf{z} \rangle|] \leq S_1 + S_2 \leq 56t^\circ = \frac{t}{64},$$

which implies that Assumption 2(a) is satisfied.

On Assumption 2(b). For $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, we have that

$$\begin{aligned} \|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} & \leq \mu_\lambda \cdot \|\langle \mathbf{a}, \mathbf{x} - \mathbf{x}' \rangle\|_{\psi_2} + \|\psi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle) - \psi_{t^\circ}(\langle \mathbf{a}, \mathbf{x}' \rangle)\|_{\psi_2} \\ & \lesssim \mu_\lambda \cdot \|\mathbf{x} - \mathbf{x}'\|_2 + \frac{\lambda}{t^\circ} \cdot \|\mathbf{x} - \mathbf{x}'\|_2, \end{aligned}$$

where we have used that ψ_{t° is $\frac{\lambda}{t^\circ}$ -Lipschitz. Since $\mu_\lambda \geq \frac{1}{2}$, $t \leq \lambda$, and $t^\circ = \frac{t}{56 \cdot 64}$, it follows that

$$\|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} \lesssim \frac{\lambda}{t} \cdot d_T(\mathbf{x}, \mathbf{x}').$$

Let $\mathbf{x} \in \mathcal{X} \subset \mathbb{S}^{p-1}$. Since $\mu_\lambda \leq 1$ and $|\psi_{t^\circ}(s)| \leq |s|$ for all $s \in \mathbb{R}$, we observe that

$$\|\xi_t(\mathbf{x})\|_{\psi_2} \leq \mu_\lambda \cdot \|\langle \mathbf{a}, \mathbf{x} \rangle\|_{\psi_2} + \|\psi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle)\|_{\psi_2} \lesssim 1.$$

Therefore, Assumption 2(b) is satisfied with $L_t \lesssim \frac{\lambda}{t}$ and $r \lesssim 1$.
 On Assumption 2(c). For $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$,

$$\|\varepsilon_t(\mathbf{x}) - \varepsilon_t(\mathbf{x}')\|_{\psi_2} = \|\phi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle) - \phi_{t^\circ}(\langle \mathbf{a}, \mathbf{x}' \rangle)\|_{\psi_2} \lesssim \frac{\lambda}{t^\circ} \cdot \|\mathbf{x} - \mathbf{x}'\|_2,$$

where we have used that ϕ_{t° is $\frac{\lambda}{t^\circ}$ -Lipschitz. Since $\mu_\lambda \geq \frac{1}{2}$ and $t^\circ = \frac{t}{56 \cdot 64}$, it follows that

$$\|\varepsilon_t(\mathbf{x}) - \varepsilon_t(\mathbf{x}')\|_{\psi_2} \lesssim \frac{\lambda}{t} \cdot d_T(\mathbf{x}, \mathbf{x}').$$

Let $\mathbf{x} \in \mathcal{X} \subset \mathbb{S}^{p-1}$. Since $|\phi_{t^\circ}(s)| \leq |s|$ for all $s \in \mathbb{R}$, we observe that

$$\|\varepsilon_t(\mathbf{x})\|_{\psi_2} = \|\phi_{t^\circ}(\langle \mathbf{a}, \mathbf{x} \rangle)\|_{\psi_2} \lesssim 1.$$

Therefore, Assumption 2(c) is satisfied with $\hat{L}_t \lesssim \frac{\lambda}{t}$ and $\hat{r} \lesssim 1$. □

Proof of Corollary 6 We follow the proof template from the beginning of Sect. 8:

Step 1a. The model setup of Corollary 6 fits into Assumption 1 as follows:

- The measurement vector $\mathbf{a} \in \mathbb{R}^p$ is centered, isotropic, and it has independent, symmetric, and sub-Gaussian coordinates with $\max_{j \in [p]} \|a_j\|_{\psi_2} \leq L$; in particular, we have that $\|\mathbf{a}\|_{\psi_2} \lesssim L$ (e.g., see [64, Lem. 3.4.2]). The signal set \mathcal{X} is a subset of $R\mathbb{B}_2^p$. The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(\mathbf{a}, \mathbf{x}) := f(\mathbf{a} \circ \mathbf{x})$, where $f: \mathbb{R}^p \rightarrow \mathbb{R}$ is given by $f(\mathbf{z}) := \sum_{j=1}^p f_j(z_j)$, with $f_j: \mathbb{R} \rightarrow \mathbb{R}$ odd, γ -Lipschitz, and satisfying the conditions (4.10) and (4.11).
- The target function $T: \mathcal{X} \rightarrow \mathbb{R}^p$ is defined by $T\mathbf{x} := \mathbb{E}[f(\mathbf{a} \circ \mathbf{x})\mathbf{a}]$.

Step 1b. The target mismatch $\rho(\mathbf{x})$ vanishes for every $\mathbf{x} \in \mathcal{X}$.

Step 1c. For the trivial choice $\tilde{y}_t(\mathbf{x}) := \tilde{y}(\mathbf{x})$, the conditions of Assumption 2 are fulfilled with

$$L_t \lesssim L \cdot \frac{\gamma}{\alpha}, \quad \hat{L}_t = 0, \quad r \lesssim L(\beta_2 + \gamma)R, \quad \hat{r} = 0.$$

Step 2. We begin by showing that

$$(T\mathbf{x})_j = \mathbb{E}[f_j(a_j x_j) a_j], \quad \text{for all } j \in [p],$$

where a_j is the j -th coordinate of the measurement vector $\mathbf{a} \in \mathbb{R}^p$. Indeed, we have that

$$\begin{aligned} (T\mathbf{x})_j &= (\mathbb{E}[f(\mathbf{a} \circ \mathbf{x})\mathbf{a}])_j = \mathbb{E}[f(\mathbf{a} \circ \mathbf{x})a_j] \\ &= \sum_{j'=1}^p \mathbb{E}[f_{j'}(a_{j'} x_{j'}) a_j] = \mathbb{E}[f_j(a_j x_j) a_j], \end{aligned}$$

where for the last equality we have used the independence of the random variables a_1, \dots, a_p and $\mathbb{E}[a_j] = 0$.

Next, let us show that $(T\mathbf{x})_j \in [\beta_1 x_j, \beta_2 x_j]$ for $x_j \geq 0$ and $(T\mathbf{x})_j \in [\beta_2 x_j, \beta_1 x_j]$ for $x_j < 0$. Let us first assume that $x_j \geq 0$. Then, clearly $\mathbb{E}[f_j(a_j x_j) a_j \chi_{(0, \infty)}(a_j)] \leq \beta_2 x_j \mathbb{E}[a_j^2 \chi_{(0, \infty)}(a_j)]$. Furthermore,

$$\begin{aligned} \mathbb{E}[f_j(a_j x_j) a_j \chi_{(-\infty, 0)}(a_j)] &= \mathbb{E}[f_j(-a_j x_j) (-a_j) \chi_{(-\infty, 0)}(a_j)] \\ &\leq \beta_2 x_j \mathbb{E}[a_j^2 \chi_{(-\infty, 0)}(a_j)], \end{aligned}$$

where we have used that f_j is an odd function for the first equality. Therefore,

$$\begin{aligned} \mathbb{E}[f_j(a_j x_j) a_j] &= \mathbb{E}[f_j(a_j x_j) a_j \chi_{(0, \infty)}(a_j)] + \mathbb{E}[f_j(a_j x_j) a_j \chi_{(-\infty, 0)}(a_j)] \\ &\leq \beta_2 x_j \mathbb{E}[a_j^2] = \beta_2 x_j, \end{aligned}$$

where $\mathbb{E}[a_j^2] = 1$ follows from the assumption that \mathbf{a} is isotropic. Similarly, we observe that

$$\mathbb{E}[f_j(a_j x_j) a_j \chi_{(0, \infty)}(a_j)] \geq \beta_1 x_j \mathbb{E}[a_j^2 \chi_{(0, \infty)}(a_j)]$$

and

$$\begin{aligned} \mathbb{E}[f_j(a_j x_j) a_j \chi_{(-\infty, 0)}(a_j)] &= \mathbb{E}[f_j(-a_j x_j) (-a_j) \chi_{(-\infty, 0)}(a_j)] \\ &\geq \beta_1 x_j \mathbb{E}[a_j^2 \chi_{(-\infty, 0)}(a_j)], \end{aligned}$$

which implies

$$\mathbb{E}[f_j(a_j x_j) a_j] \geq \beta_1 x_j \mathbb{E}[a_j^2] = \beta_1 x_j.$$

This shows that $\beta_1 x_j \leq (T\mathbf{x})_j \leq \beta_2 x_j$ for $x_j \geq 0$. Analogously, one can show that $\beta_2 x_j \leq (T\mathbf{x})_j \leq \beta_1 x_j$ for $x_j < 0$. Therefore, $|(T\mathbf{x})_j| \leq \beta_2 |x_j|$, which implies $\|T\mathbf{x}\|_2 \leq \beta_2 \|\mathbf{x}\|_2$ for every $\mathbf{x} \in \mathcal{X}$. As a consequence, it holds that $T\mathcal{X} \subset \beta_2 R \mathbb{B}_2^p$. Setting $m_0 := 0$, $\Delta := 1$, and $u_0 := u$, the claim of Corollary 6 follows directly from Theorem 2.

Step 3.

By definition of the target function it follows that $\rho(\mathbf{x}) = 0$ for every $\mathbf{x} \in \mathcal{X}$.

Step 4. We simply set $\tilde{y}_t(\mathbf{x}) := \tilde{y}(\mathbf{x})$. Then, $\varepsilon_t(\mathbf{x}) = 0$, implying that Assumption 2(a) and (c) are trivially fulfilled with $\hat{L}_t = \hat{r} = 0$. Since $\xi_t(\mathbf{x}) = \langle \mathbf{a}, T\mathbf{x} \rangle - f(\mathbf{a} \circ \mathbf{x})$, the following holds for every $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$:

$$\|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} \leq \|\langle \mathbf{a}, T\mathbf{x} - T\mathbf{x}' \rangle\|_{\psi_2} + \|f(\mathbf{a} \circ \mathbf{x}) - f(\mathbf{a} \circ \mathbf{x}')\|_{\psi_2}.$$

Clearly, we have that $\|\langle \mathbf{a}, T\mathbf{x} - T\mathbf{x}' \rangle\|_{\psi_2} \lesssim L \cdot d_T(\mathbf{x}, \mathbf{x}')$ and

$$\|f(\mathbf{a} \circ \mathbf{x}) - f(\mathbf{a} \circ \mathbf{x}')\|_{\psi_2} = \left\| \sum_{j=1}^p f_j(a_j x_j) - f_j(a_j x'_j) \right\|_{\psi_2}.$$

Let $j \in [p]$. Using that f_j is γ -Lipschitz, we obtain

$$\|f_j(a_j x_j) - f_j(a_j x'_j)\|_{\psi_2} \lesssim \gamma \cdot \|a_j x_j - a_j x'_j\|_{\psi_2} \leq L \cdot \gamma \cdot |x_j - x'_j|.$$

Furthermore, since a_j is a symmetric random variable and f_j is odd, it holds $\mathbb{E}[f_j(a_j x)] = 0$ for every $x \in \mathbb{R}$. In particular, the random variable $f_j(a_j x_j) - f_j(a_j x'_j)$ is centered. Since the random variables a_1, \dots, a_p are independent, Hoeffding’s inequality yields

$$\begin{aligned} \left\| \sum_{j=1}^p f_j(a_j x_j) - f_j(a_j x'_j) \right\|_{\psi_2}^2 &\lesssim \sum_{j=1}^p \|f_j(a_j x_j) - f_j(a_j x'_j)\|_{\psi_2}^2 \\ &\lesssim L^2 \cdot \gamma^2 \cdot \|\mathbf{x} - \mathbf{x}'\|_2^2. \end{aligned}$$

Therefore,

$$\|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} \lesssim L \cdot d_T(\mathbf{x}, \mathbf{x}') + L \cdot \gamma \cdot \|\mathbf{x} - \mathbf{x}'\|_2.$$

We now show that $d_T(\mathbf{x}, \mathbf{x}') \geq \alpha \|\mathbf{x} - \mathbf{x}'\|_2$. For $j \in [p]$, we have that

$$\begin{aligned} |(T\mathbf{x})_j - (T\mathbf{x}')_j| &= |\mathbb{E}[(f_j(a_j x_j) - f_j(a_j x'_j))a_j]| \\ &= |\mathbb{E}[(f_j(a_j x_j) - f_j(a_j x'_j))a_j \chi_{[0,\infty)}(a_j)] \\ &\quad + \mathbb{E}[(f_j(a_j x_j) - f_j(a_j x'_j))a_j \chi_{(-\infty,0)}(a_j)]|. \end{aligned}$$

We may assume that $x_j \geq x'_j$. If $a_j \geq 0$, then $a_j x_j \geq a_j x'_j$, which implies $f_j(a_j x_j) - f_j(a_j x'_j) \geq \alpha(a_j x_j - a_j x'_j)$. Therefore, $(f_j(a_j x_j) - f_j(a_j x'_j))a_j \geq \alpha(a_j x_j - a_j x'_j)a_j$, which implies

$$\begin{aligned} \mathbb{E}[(f_j(a_j x_j) - f_j(a_j x'_j))a_j \chi_{[0,\infty)}(a_j)] &\geq \mathbb{E}[\alpha(a_j x_j - a_j x'_j)a_j \chi_{[0,\infty)}(a_j)] \\ &= \alpha(x_j - x'_j)\mathbb{E}[a_j^2 \chi_{[0,\infty)}(a_j)] \geq 0. \end{aligned}$$

If $a_j < 0$, then $a_j x_j \leq a_j x'_j$, which implies $f_j(a_j x'_j) - f_j(a_j x_j) \geq \alpha(a_j x'_j - a_j x_j)$. Therefore, $f_j(a_j x_j) - f_j(a_j x'_j) \leq \alpha(a_j x_j - a_j x'_j)$, which implies $(f_j(a_j x_j) - f_j(a_j x'_j))a_j \geq \alpha(a_j x_j - a_j x'_j)a_j$. This shows

$$\begin{aligned} \mathbb{E}[(f_j(a_j x_j) - f_j(a_j x'_j))a_j \chi_{(-\infty, 0)}(a_j)] &\geq \mathbb{E}[\alpha(a_j x_j - a_j x'_j)a_j \chi_{(-\infty, 0)}(a_j)] \\ &= \alpha(x_j - x'_j)\mathbb{E}[a_j^2 \chi_{(-\infty, 0)}(a_j)] \geq 0. \end{aligned}$$

It follows that $|(T\mathbf{x})_j - (T\mathbf{x}')_j| \geq \alpha(x_j - x'_j) = \alpha|x_j - x'_j|$, and therefore $d_T(\mathbf{x}, \mathbf{x}') \geq \alpha\|\mathbf{x} - \mathbf{x}'\|_2$. Due to this inequality, we obtain

$$\begin{aligned} \|\xi_t(\mathbf{x}) - \xi_t(\mathbf{x}')\|_{\psi_2} &\lesssim L \cdot d_T(\mathbf{x}, \mathbf{x}') + L \cdot \gamma \cdot \|\mathbf{x} - \mathbf{x}'\|_2 \\ &\leq L \cdot d_T(\mathbf{x}, \mathbf{x}') + L \cdot \frac{\gamma}{\alpha} \cdot d_T(\mathbf{x}, \mathbf{x}') \\ &\leq 2 \cdot L \cdot \frac{\gamma}{\alpha} \cdot d_T(\mathbf{x}, \mathbf{x}'), \end{aligned}$$

where the last inequality follows from $\alpha \leq \gamma$.

Now, let $\mathbf{x} \in \mathcal{X}$. It holds that

$$\begin{aligned} \|\xi_t(\mathbf{x})\|_{\psi_2} &= \|\langle \mathbf{a}, T\mathbf{x} \rangle - f(\mathbf{a} \circ \mathbf{x})\|_{\psi_2} \leq \|\langle \mathbf{a}, T\mathbf{x} \rangle\|_{\psi_2} + \|f(\mathbf{a} \circ \mathbf{x})\|_{\psi_2} \\ &\lesssim L \cdot \|T\mathbf{x}\|_2 + \|f(\mathbf{a} \circ \mathbf{x})\|_{\psi_2} \leq L \cdot \beta_2 \cdot R + \|f(\mathbf{a} \circ \mathbf{x})\|_{\psi_2}. \end{aligned}$$

Since f_j is γ -Lipschitz and odd for every $j \in [p]$, we obtain

$$\|f_j(a_j x_j)\|_{\psi_2} = \|f_j(a_j x_j) - f_j(0)\|_{\psi_2} \lesssim \gamma \cdot \|a_j x_j\|_{\psi_2} \leq L \cdot \gamma \cdot |x_j|.$$

Using Hoeffding’s inequality, we conclude that

$$\|f(\mathbf{a} \circ \mathbf{x})\|_{\psi_2} = \left\| \sum_{j=1}^p f_j(a_j x_j) \right\|_{\psi_2} \lesssim \left(\sum_{j=1}^p \|f_j(a_j x_j)\|_{\psi_2}^2 \right)^{1/2} \leq L \cdot \gamma \cdot \|\mathbf{x}\|_2.$$

Hence, Assumption 2(b) is satisfied with $L_t \lesssim L \cdot \frac{\gamma}{\alpha}$ and $r \lesssim L(\beta_2 + \gamma)R$.

□

8.5 Proofs for Sect. 4.5

Proof of Corollary 7 We follow the proof template from the beginning of Sect. 8:

Step 1a. The model setup of Corollary 7 fits into Assumption 1 as follows:

- The measurement vector $\mathbf{a} \in \mathbb{R}^p$ is centered, isotropic, and has independent sub-Gaussian coordinates with $\max_{j \in [p]} \|a_j\|_{\psi_2} \leq L$; in particular, we have that $\|\mathbf{a}\|_{\psi_2} \lesssim L$ (e.g., see [64, Lem. 3.4.2]). The signal set \mathcal{X} is an arbitrary subset of $\{\mathcal{S} \subset [p] : |\mathcal{S}| \leq s\}$. The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(\mathbf{a}, \mathcal{S}) := f(\mathbf{a}_{\mathcal{S}})$, where $(\mathbf{a}_{\mathcal{S}})_j = a_j$ for $j \in \mathcal{S}$ and $(\mathbf{a}_{\mathcal{S}})_j = 0$ for $j \in \mathcal{S}^c$.

– The target function $T : \mathcal{X} \rightarrow K$ is defined by $T\mathcal{S} := \mathbb{E}[f(\mathbf{a}_{\mathcal{S}})\mathbf{a}]$.

Step 1b. The target mismatch $\rho(\mathcal{S})$ vanishes for every $\mathcal{S} \in \mathcal{X}$.

Step 1c. For the trivial choice $\tilde{y}_t(\mathcal{S}) := \tilde{y}(\mathcal{S})$, the conditions of Assumption 2 are fulfilled with

$$L_t \leq L + \gamma\alpha^{-1}, \quad \hat{L}_t = 0, \quad r \leq L\beta + \kappa, \quad \hat{r} = 0.$$

Step 2. Setting $m_0 := 0$, $\Delta := L^{-1}\sqrt{\log L}$, and $u := u_0$, the claim of Corollary 7 follows directly from Theorem 2.

Step 3. This is clear by the definition of the target function.

Step 4. We simply set $\tilde{y}_t(\mathcal{S}) := \tilde{y}(\mathcal{S})$. Then, $\varepsilon_t(\mathcal{S}) = 0$, implying that Assumption 2(a) and (c) are trivially fulfilled with $\hat{L}_t = \hat{r} = 0$. To see that Assumption 2(b) holds as well, we first recall that the coordinates of \mathbf{a} are centered and independent, so that

$$(T\mathcal{S})_j = \mathbb{E}[f(\mathbf{a}_{\mathcal{S}})a_j] = \mathbb{E}[f(\mathbf{a}_{\mathcal{S}})] \cdot \mathbb{E}[a_j] = 0 \quad \text{for all } \mathcal{S} \in \mathcal{X} \text{ and } j \in \mathcal{S}^c.$$

Together with the lower bound in (4.14), it follows that $\text{supp}(T\mathcal{S}) = \mathcal{S}$ for all $\mathcal{S} \in \mathcal{X}$. Using the lower bound in (4.14) again, we obtain the following estimate for all $\mathcal{S}, \mathcal{S}' \in \mathcal{X}$:

$$d_T(\mathcal{S}, \mathcal{S}') = \|T\mathcal{S} - T\mathcal{S}'\|_2 \geq \left(\sum_{j \in \mathcal{S} \setminus \mathcal{S}'} (T\mathcal{S})_j^2 + \sum_{j \in \mathcal{S}' \setminus \mathcal{S}} (T\mathcal{S}')_j^2 \right)^{1/2} \geq \frac{\alpha}{\sqrt{s}} \sqrt{|\mathcal{S} \Delta \mathcal{S}'|}.$$

Combining this with the assumption (4.15), it follows that

$$\begin{aligned} \|\xi_t(\mathcal{S}) - \xi_t(\mathcal{S}')\|_{\psi_2} &\leq \|\langle \mathbf{a}, T\mathcal{S} - T\mathcal{S}' \rangle\|_{\psi_2} + \|f(\mathbf{a}_{\mathcal{S}}) - f(\mathbf{a}_{\mathcal{S}'})\|_{\psi_2} \\ &\leq L \cdot d_T(\mathcal{S}, \mathcal{S}') + \gamma\alpha^{-1} \cdot d_T(\mathcal{S}, \mathcal{S}') = (L + \gamma\alpha^{-1}) \cdot d_T(\mathcal{S}, \mathcal{S}'). \end{aligned}$$

This implies $L_t \leq L + \gamma\alpha^{-1}$. Since $\text{supp}(T\mathcal{S}) = \mathcal{S}$ and $|\mathcal{S}| \leq s$ for $\mathcal{S} \in \mathcal{X}$, we also have that $|\text{supp}(T\mathcal{S})| \leq s$. The upper bound in (4.14) therefore yields $\|T\mathcal{S}\|_2 \leq \beta$ for every $\mathcal{S} \in \mathcal{X}$. Combining this estimate with (4.15), we obtain the following upper bound for the sub-Gaussian norm of $\xi_t(\mathcal{S})$:

$$\|\xi_t(\mathcal{S})\|_{\psi_2} = \|\langle \mathbf{a}, T\mathcal{S} \rangle - f(\mathbf{a}_{\mathcal{S}})\|_{\psi_2} \leq L\|T\mathcal{S}\|_2 + \|f(\mathbf{a}_{\mathcal{S}})\|_{\psi_2} \leq L\beta + \kappa.$$

Hence, $r := L\beta + \kappa$ is a valid choice for Assumption 2(b).

□

9 Proofs for Sect. 5

Proof of Theorem 3 Let \mathcal{X}_ε be a minimal subset of \mathcal{X} such that $T\mathcal{X}_\varepsilon$ is a minimal ε -net for $T\mathcal{X}$; in particular, we have that $|\mathcal{X}_\varepsilon| = |T\mathcal{X}_\varepsilon| = \mathcal{N}(T\mathcal{X}, \varepsilon)$. The assumptions of (5.2) and (5.3) allow us to first apply Theorem 2 to every $\mathbf{x}' \in \mathcal{X}_\varepsilon$ as a singleton signal set (Assumption 2 is trivially fulfilled here), and then to take a union bound over \mathcal{X}_ε . Consequently, there exist universal constants $c, C > 0$ such that the following event \mathcal{A}_1 holds with probability at least $1 - \exp(-cu^2) - \exp(-cu_0^2)$: For every $\mathbf{x}' \in \mathcal{X}_\varepsilon$ with $\rho(\mathbf{x}') \leq \frac{t}{32}$ and every $\mathbf{y} \in \mathbb{R}^m$ such that

$$\frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}')\|_{[2m_0]} \leq \Delta t \quad \text{and} \quad \frac{1}{\sqrt{m}} \sigma_{2m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}'))_2 \leq \frac{t}{20},$$

any minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - T\mathbf{x}'\|_2 \leq t$. Note that we have applied the robustness conditions of (2.4) for $2m_0$ instead of m_0 here, which is possible due to $2m_0 \leq m$.

Let \mathcal{A}_2 denote the event of (5.1) in Assumption 3. We now show that on the event $\mathcal{A}_1 \cap \mathcal{A}_2$, which occurs with probability at least $1 - \exp(-cu^2) - \exp(-cu_0^2) - \eta$, the conclusion of Theorem 3 holds: Let $\mathbf{x} \in \mathcal{X}$ be fixed and consider any input vector $\mathbf{y} \in \mathbb{R}^m$ satisfying

$$\frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_{[2m_0]} \leq \frac{1}{2} \Delta t \quad \text{and} \quad \frac{1}{\sqrt{m}} \sigma_{m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}))_2 \leq \frac{t}{40}.$$

By the definition of \mathcal{X}_ε , there exists $\mathbf{x}' \in \mathcal{X}_\varepsilon$ such that $\|T\mathbf{x} - T\mathbf{x}'\|_2 \leq \varepsilon$. According to the event \mathcal{A}_2 , we conclude that

$$\frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}')\|_{[2m_0]} \leq \frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_{[2m_0]} + \frac{1}{\sqrt{m}} \|\tilde{\mathbf{y}}(\mathbf{x}) - \tilde{\mathbf{y}}(\mathbf{x}')\|_{[2m_0]} \leq \Delta t$$

and

$$\frac{1}{\sqrt{m}} \sigma_{2m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}'))_2 \leq \frac{1}{\sqrt{m}} \sigma_{m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}))_2 + \frac{1}{\sqrt{m}} \sigma_{m_0}(\tilde{\mathbf{y}}(\mathbf{x}) - \tilde{\mathbf{y}}(\mathbf{x}'))_2 \leq \frac{t}{20}.$$

Finally, according to the event \mathcal{A}_1 , every minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - T\mathbf{x}'\|_2 \leq t$ and therefore $\|\hat{\mathbf{z}} - T\mathbf{x}\|_2 \leq \|\hat{\mathbf{z}} - T\mathbf{x}'\|_2 + \|T\mathbf{x}' - T\mathbf{x}\|_2 \leq t + \varepsilon$. □

Proof of Corollary 8 Analogously to Corollary 2, the model setup of Corollary 8 fits into Assumption 1 as follows:

- We have that $\mathbf{a} \sim \mathbf{N}(\mathbf{0}, \mathbf{I}_p)$ and therefore $\|\mathbf{a}\|_{\psi_2} \lesssim 1$. The signal set \mathcal{X} is an arbitrary subset of \mathbb{R}^p . The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(\mathbf{a}, \mathbf{x}) := \text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle)$. In particular, the resulting observation vector of \mathbf{x} is given by $\tilde{\mathbf{y}}(\mathbf{x}) = \text{sign}(\mathbf{A}\mathbf{x})$ where $\mathbf{A} \in \mathbb{R}^{m \times p}$ is a standard Gaussian random matrix with row vectors $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$.

– The target function $T : \mathcal{X} \rightarrow K$ corresponds to the (scaled) normalization

$$T\mathbf{x} := \sqrt{\frac{2}{\pi}} \frac{\mathbf{x}}{\|\mathbf{x}\|_2}.$$

As already shown in Sect. 8.1 (Step 3), the target mismatch $\rho(\mathbf{x})$ vanishes for every $\mathbf{x} \in \mathcal{X}$. Moreover, we clearly have that $r = \sup_{\mathbf{x} \in \mathcal{X}} \|\langle \mathbf{a}, T\mathbf{x} \rangle - \tilde{\mathbf{y}}(\mathbf{x})\|_{\psi_2} \lesssim 1$. Since $\phi(\beta) := \beta \sqrt{\log(e/\beta)}$ defines a continuous and non-decreasing function on $[0, 1]$ with $\phi(1) = 1$ and $\phi(0) := 0$ (by continuous extension), we may assume without loss of generality that

$$\beta \sqrt{\log(e/\beta)} = c_0 t \in (0, 1]. \tag{9.1}$$

Now, we set

$$\Delta^2 := \frac{1}{t \sqrt{\log(e/\beta)}}, \quad m_0 := \lfloor \frac{\beta}{2} m \rfloor, \quad u_0 := \sqrt{2m\beta \log(e/\beta)},$$

implying that $u_0 \geq \sqrt{2m_0 \log(em/2m_0)}$, $m_0 \in \{0, 1, \dots, \lfloor \frac{m}{2} \rfloor\}$, and $u_0^2 \geq 2c_0 t m$. The latter inequality again implies that $u_0^2 \geq u^2$, and since $u^2 \geq C_0 \cdot \log \mathcal{N}(T\mathcal{X}, \varepsilon)$, the condition of (5.3) is fulfilled. Similarly, it is not hard to see that the condition of (5.2) follows from (5.4) for C' sufficiently large (cf. Sect. 8.1 (Step 2)).

According to Theorem 4 with $H := \sqrt{\frac{\pi}{2}} T\mathcal{X} \subset \mathbb{S}^{p-1}$, there exist universal constants $c, \bar{c}, C > 0$ (possibly slightly different from those in Theorem 4) such that if

$$m \geq C \cdot \left(\varepsilon^2 \beta^{-3} \cdot w^2([T\mathcal{X}]_\varepsilon) + \beta^{-1} \cdot \log \mathcal{N}(T\mathcal{X}, \varepsilon) \right) \tag{9.2}$$

for $\varepsilon \leq \bar{c}\beta/\sqrt{\log(e/\beta)}$, the following holds with probability at least $1 - \exp(-cm\beta)$:

$$\sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \|T\mathbf{x} - T\mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{2m} \|\text{sign}(\mathbf{A}\mathbf{x}) - \text{sign}(\mathbf{A}\mathbf{x}')\|_1 \leq \frac{\beta}{2}.$$

On this event and by the above choice of m_0 , we have that

$$\sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \|T\mathbf{x} - T\mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{\sqrt{m}} \sigma_{m_0}(\tilde{\mathbf{y}}(\mathbf{x}) - \tilde{\mathbf{y}}(\mathbf{x}'))_2 = 0$$

and

$$\sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \|T\mathbf{x} - T\mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{\sqrt{m}} \|\tilde{\mathbf{y}}(\mathbf{x}) - \tilde{\mathbf{y}}(\mathbf{x}')\|_{[2m_0]} \leq \sqrt{2\beta} = \left(\frac{2c_0 t}{\sqrt{\log(e/\beta)}} \right)^{1/2} \leq \frac{1}{2} \Delta t,$$

where the last inequality holds for $c_0 > 0$ small enough. Consequently, Assumption 3 would hold for $\eta := \exp(-cm\beta)$ if we can show that (9.2) is satisfied under the hypotheses of Corollary 8. To this end, we first note that the relationship (9.1)

implies that there exists a universal constant $c'_0 > 0$ such that $\beta \geq c'_0 t / \sqrt{\log(e/t)}$. This particularly leads to the following estimates:

$$\beta \geq c'_0 t^2 \quad \text{and} \quad \frac{\beta}{\sqrt{\log(e/\beta)}} \geq \frac{(c'_0)^2 t}{c_0 \log(e/t)}.$$

Combining the first one with (5.4), it follows that $m\beta \geq c'_0 m t^2 \gtrsim u^2$, while the second one yields

$$\varepsilon \leq \frac{c' t}{\log(e/t)} \leq \frac{\bar{c}\beta}{\sqrt{\log(e/\beta)}}$$

if c' is chosen sufficiently small. Hence, (9.2) is a consequence of (5.4) and the assumption $u^2 \geq C_0 \cdot \log \mathcal{N}(T\mathcal{X}, \varepsilon)$.

Since all assumptions of Theorem 3 are satisfied, the following holds with probability at least $1 - \exp(-cu^2)$ uniformly for every $\mathbf{x} \in \mathcal{X}$: For any input vector $\mathbf{y} \in \mathbb{R}^m$ such that

$$\frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_{[2m_0]} \leq \frac{1}{2} \Delta t \quad \text{and} \quad \frac{1}{\sqrt{m}} \sigma_{m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}))_2 \leq \frac{t}{40}, \tag{9.3}$$

every minimizer $\hat{\mathbf{z}}$ of $(P_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - T\mathbf{x}\|_2 \leq t + \varepsilon \leq 2t$. The claim of Corollary 8 now follows from the fact that any input vector $\mathbf{y} \in \{-1, 1\}^m$ given by (4.1) with $\frac{1}{2m} \|\mathbf{v}\|_1 \leq \frac{\beta}{2}$ satisfies (9.3). □

Proof of Corollary 9 Analogously to Corollary 3, the model setup of Corollary 9 fits into Assumption 1 as follows:

- The measurement vector $\mathbf{a} \in \mathbb{R}^p$ is centered, isotropic, and sub-Gaussian with $\|\mathbf{a}\|_{\psi_2} \leq L$. The signal set \mathcal{X} satisfies $\mathcal{X} \subset R\mathbb{B}_2^p$. The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(\mathbf{a}, \mathbf{x}) = \text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau)$, where τ is uniformly distributed on $[-\lambda, \lambda]$ and independent of \mathbf{a} . In particular, F is a random function. Moreover, the observation vector of \mathbf{x} is given by $\tilde{\mathbf{y}}(\mathbf{x}) = \text{sign}(\mathbf{A}\mathbf{x} + \boldsymbol{\tau})$, where $\mathbf{A} \in \mathbb{R}^{m \times p}$ denotes the sub-Gaussian random matrix with row vectors $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^p$ and $\boldsymbol{\tau} := (\tau_1, \dots, \tau_m)$.
- The target function $T: \mathcal{X} \rightarrow K$ corresponds to rescaling by a factor of λ^{-1} , i.e., $T\mathbf{x} := \lambda^{-1}\mathbf{x}$.

As already shown in Sect. 8.2 (Step 3), there exists a universal constant $\tilde{C}' \geq e$ such that if

$$\lambda \geq \tilde{C}' \cdot R \cdot L \cdot \sqrt{\log(e/t)}, \tag{9.4}$$

the target mismatch satisfies $\rho(\mathbf{x}) \leq \frac{t}{32}$ for every $\mathbf{x} \in \mathcal{X}$. Moreover, we clearly have that

$$\|\langle \mathbf{a}, \lambda^{-1}\mathbf{x} \rangle - \tilde{\mathbf{y}}(\mathbf{x})\|_{\psi_2} \lesssim \lambda^{-1} \|\langle \mathbf{a}, \mathbf{x} \rangle\|_{\psi_2} + 1 \lesssim RL\lambda^{-1} + 1$$

for every $\mathbf{x} \in \mathcal{X}$, and together with (9.4), it follows that $r = \sup_{\mathbf{x} \in \mathcal{X}} \|\langle \mathbf{a}, \lambda^{-1} \mathbf{x} \rangle - \tilde{y}(\mathbf{x})\|_{\psi_2} \lesssim 1$. Since $\phi(\beta) := \beta \sqrt{\log(e/\beta)}$ defines a continuous and non-decreasing function on $[0, 1]$ with $\phi(1) = 1$ and $\phi(0) := 0$ (by continuous extension), we may assume without loss of generality that

$$\beta \sqrt{\log(e/\beta)} = c_0 L^{-2} t \in (0, 1]. \tag{9.5}$$

Now, we set

$$\Delta^2 := \frac{1}{t L^2 \sqrt{\log(e/\beta)}}, \quad m_0 := \lfloor \frac{\beta}{2} m \rfloor, \quad u_0 := \sqrt{2m\beta \log(e/\beta)},$$

implying that $u_0 \geq \sqrt{2m_0 \log(em/2m_0)}$, $m_0 \in \{0, 1, \dots, \lfloor \frac{m}{2} \rfloor\}$, and $u_0^2 \geq 2c_0 m L^{-2} t$. Combining the latter inequality with (5.5) for $C' \gtrsim L^2$ sufficiently large implies that $u_0^2 \geq u^2$, and since $u^2 \geq C_0 \cdot \log \mathcal{N}(\lambda^{-1} \mathcal{X}, \varepsilon)$, the condition of (5.3) is fulfilled. Similarly, it is not hard to see that the condition of (5.2) follows from (5.5) for $C' \gtrsim L^2 \log L$ sufficiently large (cf. Sect. 8.2 (Step 2)).

According to Theorem 5, there exist constants $c, \bar{c}, C, \tilde{C} > 0$ only depending on L (possibly slightly different from those in Theorem 5) such that if $\lambda \geq \tilde{C} \cdot R$ and

$$m \geq C \cdot \left(\varepsilon^2 \beta^{-3} \cdot w^2([\!(2\lambda)^{-1} \mathcal{X}]_\varepsilon) + \beta^{-1} \cdot \log \mathcal{N}(\lambda^{-1} \mathcal{X}, \varepsilon) \right) \tag{9.6}$$

for $\varepsilon \leq \bar{c} \beta / \sqrt{\log(e/\beta)}$, the following holds with probability at least $1 - \exp(-cm\beta)$:

$$\sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \lambda^{-1} \|\mathbf{x} - \mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{2m} \|\text{sign}(\mathbf{A}\mathbf{x} + \boldsymbol{\tau}) - \text{sign}(\mathbf{A}\mathbf{x}' + \boldsymbol{\tau})\|_1 \leq \frac{\beta}{2}.$$

On this event and by the above choice of m_0 , we have that

$$\sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \lambda^{-1} \|\mathbf{x} - \mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{\sqrt{m}} \sigma_{m_0}(\tilde{\mathbf{y}}(\mathbf{x}) - \tilde{\mathbf{y}}(\mathbf{x}'))_2 = 0$$

and

$$\sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \lambda^{-1} \|\mathbf{x} - \mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{\sqrt{m}} \|\tilde{\mathbf{y}}(\mathbf{x}) - \tilde{\mathbf{y}}(\mathbf{x}')\|_{[2m_0]} \leq \sqrt{2\beta} = \left(\frac{2c_0 t}{L^2 \sqrt{\log(e/\beta)}} \right)^{1/2} \leq \frac{1}{2} \Delta t,$$

where the last inequality holds for $c_0 > 0$ small enough. Consequently, Assumption 3 would hold for $\eta := \exp(-cm\beta)$ if we can show that (9.6) is satisfied under the hypotheses of Corollary 9. To this end, we first note that the relationship (9.5) implies that there exists a universal constant $c'_0 > 0$ such that $\beta \geq c'_0 t L^{-2} / \sqrt{\log(eL^2/t)}$. This

particularly leads to the following estimates:

$$\beta \geq c'_0 t^2 L^{-2} \quad \text{and} \quad \frac{\beta}{\sqrt{\log(e/\beta)}} \geq \frac{(c'_0)^2 t}{c_0 L^2 \log(eL^2/t)}.$$

Combining the first one with (5.5), it follows that $m\beta \geq c'_0 m t^2 L^{-2} \gtrsim u^2$ for $C' \gtrsim L^2$ sufficiently large, while the second one yields

$$\varepsilon \leq \frac{c' t}{\log(e/t)} \leq \frac{\bar{c}\beta}{\sqrt{\log(e/\beta)}}$$

if c' is chosen sufficiently small (depending on L). Hence, (9.6) is a consequence of (5.5) and the assumption $u^2 \geq C_0 \cdot \log \mathcal{N}(\lambda^{-1}\mathcal{X}, \varepsilon)$.

Since all assumptions of Theorem 3 are satisfied, the following holds with probability at least $1 - \exp(-cu^2)$ uniformly for every $\mathbf{x} \in \mathcal{X}$: For any input vector $\mathbf{y} \in \mathbb{R}^m$ such that

$$\frac{1}{\sqrt{m}} \|\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x})\|_{[2m_0]} \leq \frac{1}{2} \Delta t \quad \text{and} \quad \frac{1}{\sqrt{m}} \sigma_{m_0}(\mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}))_2 \leq \frac{t}{40}, \tag{9.7}$$

every minimizer $\hat{\mathbf{z}}$ of $(\mathbf{P}_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \lambda^{-1}\mathbf{x}\|_2 \leq t + \varepsilon \leq 2t$. The claim of Corollary 9 now follows from the fact that any input vector $\mathbf{y} \in \{-1, 1\}^m$ given by (4.3) with $\frac{1}{2m} \|\mathbf{v}\|_1 \leq \frac{\beta}{2}$ satisfies (9.7). □

10 Proofs for Sect. 3

Proof of Corollary 1 We follow the proof template from the beginning of Sect. 8:

Step 1a. The model setup of Corollary 1 fits into Assumption 1 as follows:

- The measurement vector $\mathbf{a} \in \mathbb{R}^p$ is centered, isotropic, and sub-Gaussian with $\|\mathbf{a}\|_{\psi_2} \leq L$. The signal set \mathcal{X} is a bounded subset of \mathbb{R}^p . The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(\mathbf{a}, \mathbf{x}) := \langle \mathbf{a}, \mathbf{x} \rangle + \tau$, where $\tau \sim \mathcal{N}(0, \sigma^2)$.
- The target function $T: \mathcal{X} \rightarrow K$ is the canonical embedding into K , i.e., $T\mathbf{x} := \mathbf{x}$. In particular, we have that $d_T(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_2$.

Step 1b. The target mismatch $\rho(\mathbf{x})$ vanishes for every $\mathbf{x} \in \mathcal{X}$.

Step 1c. For the trivial choice $\tilde{\mathbf{y}}_t(\mathbf{x}) := \tilde{\mathbf{y}}(\mathbf{x})$, the conditions of Assumption 2 are fulfilled with

$$L_t = 0, \quad \hat{L}_t = 0, \quad r \lesssim \sigma, \quad \hat{r} = 0.$$

Step 2. Setting $m_0 := 0$, $\Delta := L^{-1} \sqrt{\log L}$, and $u := u_0$, the claim of Corollary 1 follows directly from Theorem 2.

Step 3. Let $\mathbf{x} \in \mathcal{X}$. By the isotropy of \mathbf{a} , we have that $\mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] = \mathbb{E}[\langle \mathbf{a}, \mathbf{x} \rangle \mathbf{a}] = \mathbf{x} = T\mathbf{x}$, and therefore $\rho(\mathbf{x}) = 0$.

Step 4. We simply set $\tilde{y}_t(\mathbf{x}) := \tilde{y}(\mathbf{x})$. Then, $\varepsilon_t(\mathbf{x}) = 0$, implying that Assumption 2(a) and (c) are trivially fulfilled with $\hat{L}_t = \hat{r} = 0$. Furthermore, since $\xi_t(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle - \tilde{y}_t(\mathbf{x}) = -\tau$, Assumption 2(b) holds with $L_t = 0$ and $r := \|\tau\|_{\psi_2} \lesssim \sigma$, since $\tau \sim \mathcal{N}(0, \sigma^2)$.

□

Proof of Inequality (3.3) Since $\|\mathbf{x}\|_1 = R$ for every $\mathbf{x} \in \mathcal{X}$, we may use [54, Lem. 4.5] to observe that

$$\text{cone}(R\mathbb{B}_1^p - \mathbf{x}) \cap \mathbb{S}^{p-1} \subset H := 2 \text{conv}(\mathbf{v} \in \mathbb{S}^{p-1} : \|\mathbf{v}\|_0 \leq s),$$

where $\text{conv}(\cdot)$ denotes the convex hull. Moreover, [49, Lem. 2.3] yields $w(H) \lesssim \sqrt{s \cdot \log(2p/s)}$. Therefore, we obtain

$$\begin{aligned} w_t(K - \mathcal{X}) &\leq w_0(K - \mathcal{X}) \\ &= w(\text{cone}(K - \mathcal{X}) \cap \mathbb{S}^{p-1}) \\ &= w(\bigcup_{\mathbf{x} \in \mathcal{X}} \text{cone}(R\mathbb{B}_1^p - \mathbf{x}) \cap \mathbb{S}^{p-1}) \\ &\leq w(H) \lesssim \sqrt{s \cdot \log(2p/s)}, \end{aligned}$$

as claimed in (3.3). □

Proof of Proposition 1 We begin with writing out the definition of the local mean width as follows:

$$w_t(K - \mathcal{X}) \leq \frac{1}{t} \cdot w((K - \mathcal{X}) \cap t\mathbb{B}_2^p) = \frac{1}{t} \cdot \mathbb{E} \left[\sup_{\substack{\mathbf{z} \in K, \mathbf{x} \in \mathcal{X} \\ \|\mathbf{z} - \mathbf{x}\|_2 \leq t}} \langle \mathbf{g}, \mathbf{z} - \mathbf{x} \rangle \right],$$

where $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$. Now, let \mathcal{X}_t be a minimal t -net for \mathcal{X} ; in particular, we have that $|\mathcal{X}_t| = \mathcal{N}(\mathcal{X}, t)$. Hence, for every $\mathbf{z} \in K$ and $\mathbf{x} \in \mathcal{X}$ with $\|\mathbf{z} - \mathbf{x}\|_2 \leq t$, there exists $\mathbf{v}_x \in \mathcal{X}_t$ such that $\|\mathbf{v}_x - \mathbf{x}\|_2 \leq t$ and $\|\mathbf{v}_x - \mathbf{z}\|_2 \leq 2t$. Therefore, we obtain

$$\begin{aligned} \mathbb{E} \left[\sup_{\substack{\mathbf{z} \in K, \mathbf{x} \in \mathcal{X} \\ \|\mathbf{z} - \mathbf{x}\|_2 \leq t}} \langle \mathbf{g}, \mathbf{z} - \mathbf{x} \rangle \right] &= \mathbb{E} \left[\sup_{\substack{\mathbf{z} \in K, \mathbf{x} \in \mathcal{X} \\ \|\mathbf{z} - \mathbf{x}\|_2 \leq t}} \langle \mathbf{g}, \mathbf{z} - \mathbf{v}_x \rangle + \langle \mathbf{g}, \mathbf{v}_x - \mathbf{x} \rangle \right] \\ &\leq \mathbb{E} \left[\sup_{\substack{\mathbf{z} \in K, \mathbf{v} \in \mathcal{X}_t \\ \|\mathbf{z} - \mathbf{v}\|_2 \leq 2t}} \langle \mathbf{g}, \mathbf{z} - \mathbf{v} \rangle + \sup_{\substack{\mathbf{x} \in \mathcal{X}, \mathbf{v} \in \mathcal{X}_t \\ \|\mathbf{x} - \mathbf{v}\|_2 \leq t}} \langle \mathbf{g}, \mathbf{v} - \mathbf{x} \rangle \right] \\ &= w((K - \mathcal{X}_t) \cap 2t\mathbb{B}_2^p) + w((\mathcal{X}_t - \mathcal{X}) \cap t\mathbb{B}_2^p) \\ &= w(\bigcup_{\mathbf{v} \in \mathcal{X}_t} (K - \mathbf{v}) \cap 2t\mathbb{B}_2^p) + \underbrace{w((\mathcal{X}_t - \mathcal{X}) \cap t\mathbb{B}_2^p)}_{\leq w(\mathcal{X}_t - \mathcal{X}) \leq 2 \cdot w(\mathcal{X})}. \end{aligned}$$

Next, we make use of the fact that \mathcal{X}_t is a finite set. Indeed, applying [15, Lem. 6.1], this allows us to “pull out” the union over \mathcal{X}_t in the following way:

$$w(\bigcup_{\mathbf{v} \in \mathcal{X}_t} (K - \mathbf{v}) \cap 2t\mathbb{B}_2^p) \lesssim \sup_{\mathbf{v} \in \mathcal{X}_t} w((K - \mathbf{v}) \cap 2t\mathbb{B}_2^p) + t \cdot \sqrt{\log \mathcal{N}(\mathcal{X}, t)}.$$

By Sudakov minoration (e.g., see [64, Thm. 7.4.1]), we have that $t \cdot \sqrt{\log \mathcal{N}(\mathcal{X}, t)} \lesssim w(\mathcal{X})$. Hence, overall, we have that

$$\begin{aligned} w_t(K - \mathcal{X}) &\lesssim \sup_{\mathbf{v} \in \mathcal{X}_t} \underbrace{w\left(\frac{1}{2t}(K - \mathbf{v}) \cap \mathbb{B}_2^p\right)}_{=\tilde{w}_t(K - \mathbf{v})} + t^{-1} \cdot w(\mathcal{X}) \\ &\lesssim \sup_{\mathbf{x} \in \mathcal{X}} \tilde{w}_t(K - \mathbf{x}) + t^{-1} \cdot w(\mathcal{X}). \end{aligned}$$

The “moreover”-part of Proposition 1 follows directly from a stability bound derived in [23, Prop. 2.6]. \square

Acknowledgements The authors thank Sjoerd Dirksen and Maximilian März for fruitful discussions. Moreover, the authors would like to thank the anonymous referees for their very useful comments and suggestions which have helped to improve the present article.

References

1. Ai, A., Lapanowski, A., Plan, Y., Vershynin, R.: One-bit compressed sensing with non-Gaussian measurements. *Linear Algebra Appl.* **441**, 222–239 (2014)
2. Amelunxen, D., Lotz, M., McCoy, M.B., Tropp, J.A.: Living on the edge: phase transitions in convex programs with random data. *Inf. Inference* **3**(3), 224–294 (2014)
3. Baraniuk, R.G., Foucart, S., Needell, D., Plan, Y., Wootters, M.: Exponential decay of reconstruction error from binary measurements of sparse signals. *IEEE Trans. Inf. Theory* **63**(6), 3368–3385 (2017)
4. Bhandari, A., Krahmer, F., Raskar, R.: On unlimited sampling and reconstruction. *IEEE Trans. Signal Process.* **69**, 3827–3839 (2020)
5. Bora A, Jalal A, Price E, Dimakis AG. Compressed sensing using generative models. In: D. Precup, Y.W. Teh (eds.) *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017;70:537–546
6. Boufounos PT, Jacques L, Krahmer F, Saab R. Quantization and compressive sensing. In: H. Boche, R. Calderbank, G. Kutyniok, J. Vybiral (eds.) *Compressed Sensing and its Applications: MATHEON Workshop 2013, Applied and Numerical Harmonic Analysis*, pp. 193–237. Birkhäuser Cham 2015
7. Brillinger DR. A generalized linear model with “Gaussian” regressor variables. In: P.J. Bickel, K. Doksum, J. Hodges (eds.) *A Festschrift For Erich L. Lehmann*. Chapman and Hall/CRC 1982;pp. 97–114
8. Cai, J.F., Xu, W.: Guarantees of total variation minimization for signal recovery. *Inf. Inference* **4**(4), 328–353 (2015)
9. Candès, E.J., Eldar, Y.C., Strohmer, T., Vershynin, V.: Phase retrieval via matrix completion. *SIAM J. Imaging Sci.* **6**(1), 199–225 (2013)
10. Candès, E.J., Romberg, J.K., Tao, T.: Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory* **52**(2), 489–509 (2006)
11. Candès, E.J., Romberg, J.K., Tao, T.: Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.* **59**(8), 1207–1223 (2006)
12. Chandrasekaran, V., Recht, B., Parrilo, P.A., Willsky, A.S.: The convex geometry of linear inverse problems. *Found. Comput. Math.* **12**(6), 805–849 (2012)
13. Dabeer, O., Karnik, A.: Signal parameter estimation using 1-bit dithered quantization. *IEEE Trans. Inf. Theory* **52**(12), 5389–5405 (2006)
14. Dirksen S. Quantized compressed sensing: A survey. In: H. Boche, G. Caire, R. Calderbank, G. Kutyniok, R. Mathar, P. Petersen (eds.) *Compressed Sensing and Its Applications: Third International MATHEON Conference 2017, Applied and Numerical Harmonic Analysis*, . Birkhäuser Cham 2019;67–95
15. Dirksen S, Genzel M, Jacques L, Stollenwerk A. The separation capacity of random neural networks 2021. Preprint 2108.00207
16. Dirksen, S., Jung, H.C., Rauhut, H.: One-bit compressed sensing with partial Gaussian circulant matrices. *Inf. Inference* **9**(3), 601–626 (2020)

17. Dirksen S, Mendelson S. Non-Gaussian hyperplane tessellations and robust one-bit compressed sensing 2018. Preprint 1805.09409
18. Dirksen S, Mendelson S. Robust one-bit compressed sensing with partial circulant matrices 2018. Preprint 1812.06719
19. Donoho, D.L.: Compressed sensing. *IEEE Trans. Inf. Theory* **52**(4), 1289–1306 (2006)
20. Foucart S. Flavors of compressive sensing. In: G.E. Fasshauer, L.L. Schumaker (eds.) *Approximation Theory XV: San Antonio 2016*, Springer Proceedings in Mathematics & Statistics, 2017;61–104
21. Foucart S, Rauhut H. *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis. Birkhäuser Basel 2013
22. Genzel, M.: High-dimensional estimation of structured signals from non-linear observations with general convex loss functions. *IEEE Trans. Inf. Theory* **63**(3), 1601–1619 (2017)
23. Genzel M. The mismatch principle and ℓ^1 -analysis compressed sensing: A unified approach to estimation under large model uncertainties and structural constraints. Phd thesis, available online: <https://doi.org/10.14279/depositonce-8394>, Technische Universität Berlin 2019
24. Genzel, M., Jung, P.: Recovering structured data from superimposed non-linear measurements. *IEEE Trans. Inf. Theory* **66**(1), 453–477 (2020)
25. Genzel M, Kipp C. Generic error bounds for the generalized Lasso with sub-exponential data 2020. Preprint 2004.05361
26. Genzel, M., Kutyniok, G., März, M.: ℓ^1 -analysis minimization and generalized (co-)sparsity: when does recovery succeed?. *Appl. Comput. Harmon. Anal.* **52**, 82–140 (2021)
27. Genzel M, März M, Seidel R. (2021) Compressed sensing with 1D total variation: breaking sample complexity barriers via non-uniform recovery. *Inf Inference*. <https://doi.org/10.1093/imaiai/iaab001>
28. Genzel, M., Stollenwerk, A.: Robust 1-bit compressed sensing via hinge loss minimization. *Inf. Inference* **9**(2), 361–422 (2020)
29. Goldstein, L., Minsker, S., Wei, X.: Structured signal recovery from non-linear and heavy-tailed measurements. *IEEE Trans. Inf. Theory* **64**(8), 5513–5530 (2018)
30. Gordon, Y.: On Milman’s inequality and random subspaces which escape through a mesh in \mathbb{R}^n . In: J. Lindenstrauss, V.D. Milman (eds.) *Geometric Aspects of Functional Analysis*, Springer Berlin Heidelberg, New York (1988)
31. Gray, R.M., Neuhoff, D.L.: Quantization. *IEEE Trans. Inf. Theory* **44**(6), 2325–2383 (1998)
32. Gray, R.M., Stockham, T.G.: Dithered quantizers. *IEEE Trans. Inf. Theory* **39**(3), 805–812 (1993)
33. Jacques, L., Cambareri, V.: Time for dithering: fast and quantized random embeddings via the restricted isometry property. *Inf. Inference* **6**(4), 441–476 (2017)
34. Jacques, L., Laska, J.N., Boufounos, P.T., Baraniuk, R.G.: Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Trans. Inf. Theory* **59**(4), 2082–2102 (2013)
35. Jeong H, Li X, Plan Y, Yılmaz Ö. Sub-gaussian matrices on sets: Optimal tail dependence and applications 2020. Preprint 2001.10631
36. Jung HC, Maly J, Palzer L, Stollenwerk A. Quantized compressed sensing by rectified linear units 2019. Preprint 1911.07816
37. Knudson, K., Saab, R., Ward, R.: One-bit compressive sensing with norm estimation. *IEEE Trans. Inf. Theory* **62**(5), 2748–2758 (2016)
38. Krahmer F, Kruschel C, Sandbichler M. Total Variation Minimization in Compressed Sensing. In: H. Boche, G. Caire, R. Calderbank, M. März, G. Kutyniok, R. Mathar (eds.) *Compressed Sensing and its Applications: Second International MATHEON Conference 2015.* Birkhäuser 2017;pp. 333–358
39. Liu Z, Scarlett J. The generalized lasso with nonlinear observations and generative priors 2020. Preprint 2006.12415
40. März M, Boyer C, Kahn J, Weiss P. Sampling rates for ℓ^1 -synthesis 2020. Preprint 2004.07175
41. Mendelson S. Learning without concentration. *J. ACM* **62**(3): 21 (2015)
42. Mendelson, S.: Upper bounds on product and multiplier empirical processes. *Stoch. Proc. Appl.* **126**(12), 3652–3680 (2016)
43. Mendelson, S., Pajor, A., Tomczak-Jaegermann, N.: Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geom. Funct. Anal.* **17**(4), 1248–1282 (2007)
44. Moshtaghpour, A., Jacques, L., Cambareri, V., Degraux, K., De Vleeschouwer, C.: Consistent basis pursuit for signal and matrix estimates in quantized compressed sensing. *IEEE Signal Process. Lett.* **23**(1), 25–29 (2016)
45. Oymak S, Recht B. Near-optimal bounds for binary embeddings of arbitrary sets 2015. Preprint 1512.04433

46. Oymak, S., Recht, B., Soltanolkotabi, M.: Sharp time-data tradeoffs for linear inverse problems. *IEEE Trans. Inf. Theory* **64**(6), 4129–4158 (2018)
47. Oymak, S., Soltanolkotabi, M.: Fast and reliable parameter estimation from nonlinear observations. *SIAM J. Optim.* **27**(4), 2276–2300 (2017)
48. Plan, Y., Vershynin, R.: One-bit compressed sensing by linear programming. *Comm. Pure Appl. Math.* **66**(8), 1275–1297 (2013)
49. Plan, Y., Vershynin, R.: Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach. *IEEE Trans. Inf. Theory* **59**(1), 482–494 (2013)
50. Plan, Y., Vershynin, R.: Dimension reduction by random hyperplane tessellations. *Discrete Comput. Geom.* **51**(2), 438–461 (2014)
51. Plan, Y., Vershynin, R.: The generalized Lasso with non-linear observations. *IEEE Trans. Inf. Theory* **62**(3), 1528–1537 (2016)
52. Plan, Y., Vershynin, R., Yudovina, E.: High-dimensional estimation with geometric constraints. *Inf. Inference* **6**(1), 1–40 (2016)
53. Richard, E., Obozinski, G.R., Vert, J.P.: Tight convex relaxations for sparse matrix factorization. *Advances in Neural Information Processing Systems 27*: 3284–3292 (2014)
54. Rudelson, M., Vershynin, R.: On sparse reconstruction from fourier and gaussian measurements. *Comm. Pure Appl. Math.* **61**(8), 1025–1045 (2008)
55. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* **60**(1–4), 259–268 (1992)
56. Sattar, Y., Oymak, S.: Quickly finding the best linear model in high dimensions via projected gradient descent. *IEEE Trans. Signal Process.* **68**, 818–829 (2020)
57. Stojnic M. Various thresholds for ℓ_1 -optimization in compressed sensing 2009. Preprint 0907.3666
58. Stollenwerk A. One-bit compressed sensing and fast binary embeddings. Phd thesis, available online: <https://doi.org/10.18154/RWTH-2020-00296>, RWTH Aachen 2019
59. Talagrand, M.: Upper and Lower Bounds for Stochastic Processes, *Ergebnisse der Mathematik und ihrer Grenzgebiete*, vol. 3. Springer Berlin Heidelberg (2014)
60. Thrampoulidis C, Abbasi E, Hassibi B. LASSO with non-linear measurements is equivalent to one with linear measurements. In: C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, R. Garnett (eds.) *Advances in Neural Information Processing Systems 28*, pp. 3402–3410. Curran Associates 2015
61. Thrampoulidis C, Rawat AS. Lifting high-dimensional non-linear models with Gaussian regressors 2017. Preprint 1712.03638
62. Thrampoulidis, C., Rawat, A.S.: The generalized Lasso for sub-gaussian measurements with dithered quantization. *IEEE Trans. Inf. Theory* **66**(4), 2487–2500 (2020)
63. Tropp, J.A.: Convex recovery of a structured signal from independent random linear measurements. In: G.E. Pfander (ed.) *Sampling Theory, a Renaissance, Applied and Numerical Harmonic Analysis*, pp. 76–101. Birkhäuser Cham (2015)
64. Vershynin R. High-Dimensional Probability: An Introduction with Applications in Data Science, *Cambridge Series in Statistical and Probabilistic Mathematics*, vol. 47. Cambridge University Press 2018
65. Xu, C., Jacques, L.: Quantized compressive sensing with RIP matrices: the benefit of dithering. *Inf. Inference* **9**(3), 543–586 (2020)
66. Yang Z, Balasubramanian K, Wang Z, Liu H. Estimating high-dimensional non-Gaussian multiple index models via Stein’s lemma. In: I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (eds.) *Advances in Neural Information Processing Systems 30*, pp. 6097–6106. Curran Associates 2017

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.