

A Unified Approach to Uniform Signal Recovery From Non-Linear Observations

Martin Genzel*

Alexander Stollenwerk†

Abstract. Recent advances in quantized compressed sensing and high-dimensional estimation have shown that signal recovery is even feasible under strong non-linear 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 non-linear observations is still missing. This paper develops a unified approach to this problem under the assumption 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 non-linearity. Based on empirical process theory, a key technical novelty is an approximative increment condition that can be implemented for all common types of non-linear models. This flexibility allows us to apply our approach to a variety of problems in quantized 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 can be incorporated into our framework in a plug-and-play manner, thereby implying near-optimal error bounds.

Key words. Uniform recovery, high-dimensional estimation, non-linear observations, quantized compressed sensing, empirical processes.

1 Introduction

This paper is concerned with the following fundamental signal recovery task:

Problem 1.1 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} can be 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 1.1 is the approach of *compressed sensing*. Dating back to the seminal works of Candès, Romberg, Tao, and Donoho [CRT06a; CRT06b; Don06], 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 very regime, Problem 1.1 is nowadays fairly well understood, corroborated by various real-world applications, a wide range of algorithmic methods, and a rich theoretical foundation; see [FR13] for a comprehensive overview. Put in very simple terms, the success of compressed sensing is based on the fact 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 one is wondering whether it can be used to recover *all* (or most) signals in \mathcal{X} . But also apart from this practical concern, the above quest for uniform recovery is an interesting mathematical problem in its own right, as 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 the study of Problem 1.1 in the much less explored situation of *non-linear* output functions. In fact, many conceptions from compressed sensing theory, such as the

*Utrecht University, Mathematical Institute, Utrecht, Netherlands (e-mail: m.genzel@uu.nl).

†UCLouvain, ICTEAM Institute, ISPGROUP, Louvain-la-Neuve, Belgium (e-mail: alexander.stollenwerk@uclouvain.be).

restricted isometry and nullspace property, are tailored to linear models and do not carry over directly to the non-linear case. However, as we will see in the next subsection, the presence of non-linear measuring components is not merely an academic concern but relevant to many problems in signal processing and high-dimensional statistics.¹

1.1 Prior Art

There essentially 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, in practice, analog (linear) measurements often need to be quantized before further processing. A very 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 $y_i = \text{sign}(\langle a_i, x \rangle)$. Due to this considerable loss of information, it may come as a surprise that tractable recovery methods are still available and Problem 1.1 is relatively well understood in the situation of quantized measurements. In fact, a solid theoretical basis as well as efficient algorithms have been developed over the last few years, including significant progress on lower bounds [DM18b; JLB13], robustness [DM18a; PV13b], advanced quantization schemes [BFNPW17; JMPS19; KSW16; MJCDD16; XJ18], and non-Gaussian measurements [ALPV14; DJR17; DM18a; DM18b]. While this list of references is certainly incomplete, the surveys of Dirksen [Dir19] and Boufounos et al. [BJKS15] provide a nice overview of recent advances in quantized compressed sensing as well as background material.

These achievements have also given rise to several important mathematical tools. One of the most notable breakthroughs is due to the use of quantized embedding results for signal recovery—a highly geometric argument based on uniform, random hyperplane tessellations, e.g., see [DM18a; JLB13; OR15; PV13b; PV14]. Closer to the original ideas of compressed sensing are recent reconstruction guarantees relying on variants of the restricted isometry property, e.g., see [DJR17; Fou17; JC17; XJ18]. However, these techniques are strongly tailored to quantized measurements and it remains largely unclear how they can be extended to other instances of Problem 1.1.

The second branch of related literature is much less restrictive with respect to the underlying model and allows for observations of the form $y_i = f(\langle a_i, x \rangle) + v_i$ in (1.1), where $f: \mathbb{R} \rightarrow \mathbb{R}$ can be non-linear and random. A pioneering work on these so-called *single-index models* is the one of Plan and Vershynin [PV16] (inspired by ideas of Brillinger [Bri82]), who study the *generalized Lasso* as recovery method:

$$\min_{z \in K} \frac{1}{m} \sum_{i=1}^m (y_i - \langle a_i, z \rangle)^2. \quad (\mathbf{P}_{K,y})$$

Here, $K \subset \mathbb{R}^p$ is a *convex* constraint set, which can be seen as an appropriate relaxation of the actual signal set \mathcal{X} and thereby makes $(\mathbf{P}_{K,y})$ tractable in many situations of interest. Although this simple “linearization” strategy might appear very coarse, it can be shown to produce satisfactory outcomes for Gaussian measurements, even if f is highly non-linear. A key benefit of using $(\mathbf{P}_{K,y})$ is that it does not require any (explicit) knowledge of the observation model, which enables various applications to signal processing and statistical estimation problems, e.g., see [Gen19, Chap. 3 & 4]. The work of Plan and Vershynin [PV16] is just an important example of a whole line of research, consisting of many related and follow-up works, e.g., see [Gen17; GK18; GS20; GMW18; OS17; PVY16; SO20; TAH15; TR20]; remarkably, this approach even extends to phase-retrieval-like problems where f is an even function [GK20; TR17; YBWL17]. For a more detailed discussion of the literature, we refer to [PVY16, Sec. 6] and [Gen19, Sec. 4.2].

At first sight, the aforementioned works provide a sound solution to Problem 1.1, 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 $x \in \mathcal{X}$. This

¹From a statistical perspective, Problem 1.1 describes a *semi-parametric* estimation task, and in principle, all findings of the present article can be viewed in this context (cf. [Gen19; GK20; GK18]). However, since our primary focus is on *uniform* recovery, we have decided to stick to the terminology of signal processing and compressed sensing.

strategy naturally leads to *non-uniform* recovery guarantees and we are not aware of any uniform extension in that regard. Finally, we would like to point out that the two research areas discussed above are not independent but exhibit certain overlaps, concerning both their main messages and proof strategies, for instance, see [GS20; TR20].

1.2 Contributions and Overview

The purpose of this work is to fill the striking gap between the lines of research discussed in the previous subsection and thereby, to a certain extent, to get the “best out of both worlds”. On the one hand, we show that uniform recovery is not only limited to linear or quantized measurement schemes but applies to a much larger class 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 recovery methodology suggested by the second branch of literature from Subsection 1.1, analyzing the performance of the generalized Lasso $(P_{K,y})$. Indeed, our main result, Theorem 2.7 in Subsection 2.3, establishes a uniform error bound for $(P_{K,y})$ under very general conditions, including all observation models mentioned above. In this context, it is worth pointing out that the present paper is a contribution to Problem 1.1 in the first place, but we do *not* promote the Lasso as the estimator of choice in practice.² While our findings imply that there exists a “universal” recovery strategy, which often yields very satisfactory outcomes, it might be outperformed by approaches that are specifically tailored to a certain model situation. In particular, we have decided to focus on $(P_{K,y})$ mainly because of its simplicity and popularity, but similar results could be shown for other recovery methods as well (cf. Remark 2.8(4)).

The reconstruction guarantee of Theorem 2.7 is formulated in a quite general, yet abstract setting, which requires a certain technical preparation. However, in order to familiarize the reader with the results of this paper, we now state a special case for single-index models with Lipschitz continuous output functions and Gaussian measurements (see Subsection 7.4 for a proof). Note that some of the notation from Subsection 1.3 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 2.6.

Theorem 1.2 *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 $\mathbf{a} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be γ -Lipschitz, and set $\mu := \mathbb{E}[f(g)g]$ and $r := \|f(g) - \mu g\|_{\psi_2}$ for $g \sim \mathcal{N}(0, 1)$. 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

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

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

Theorem 1.2 can be seen as a uniform version of Plan’s and Vershynin’s main result on single-index models [PV16, 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

²In particular, we will not discuss any details about possible algorithmic implementations of $(P_{K,y})$, which is an important subject in its own right.

between m , t , and the complexity parameters, we refer to the discussion of Theorem 2.7 in Subsection 2.3. To the best of our knowledge, Theorem 1.2 is a new result in its own right, and in particular, it is not a consequence of the approaches outlined in the previous subsection. In fact, a novel argument is required to prove such a type of recovery guarantee. The strategy we follow in this work is based on recent advances in empirical process theory and generic chaining. A particularly noteworthy ingredient of our proof is a powerful concentration inequality for empirical product processes due to Mendelson [Men16], which allows us to control the multiplier term of the excess loss of $(P_{K,y})$ uniformly over \mathcal{X} .

The main result of Theorem 2.7 extends Theorem 1.2 by several important aspects, being in line with the main achievements of our approach:

- *Discontinuous output functions.* The Lipschitz assumption in Theorem 1.2 prohibits popular models like quantized observations, and it turns out that an extension in that regard is not obvious. To this end, we introduce an *approximative increment condition* on the observation variable, which is satisfied for many discontinuous non-linearities; see Subsection 2.2 for more details and Subsection 3.1–3.3 for applications to quantized compressed sensing.
- *Beyond Gaussian measurements.* Compared to the previous point, it is relatively straightforward to allow for i.i.d. sub-Gaussian measurement vectors in Theorem 1.2. However, this relaxation leads to an additional constraint term, called the *target mismatch*, whose size depends on the distribution of the measurement vectors and the actual observation model; see Definition 2.2 in Subsection 2.1.
- *Beyond single-index models.* In Section 2, we will study the more general situation of (1.1), where y_i is not necessarily representable in terms of the inner product $\langle a_i, x \rangle$. While this requires an additional step of abstraction, it enables for more complicated models as well as refined recovery tasks, e.g., support estimation (see Subsection 3.5). For this purpose, we will introduce so-called *target functions* in Subsection 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 1.2, this transformation simply corresponds to rescaling x by the scalar factor μ .
- *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 1.2 is standard for linear models and also remains useful in the situation of Lipschitz non-linearities. However, such a condition is overly restrictive for quantized measurement schemes, since the noise is due to (few) wrong bits here instead of (small) real-valued perturbations. Our main result Theorem 2.7 therefore involves an important relaxation in this respect, permitting a certain set of “gross” outliers in the observation vector.

The purpose of our methodology is to enable a more systematic study of Problem 1.1 than before, especially when it comes to non-linear observation models. To a certain degree, this paper promotes an alternative path to compressed sensing theory that does without common tools like the restricted isometry property or quantized embedding results. But despite many overlapping results, each approach clearly comes along with its own strengths as well as its (in-)accessible regimes; see Section 5 for a more detailed discussion in this regard. Beyond that, we will demonstrate in Section 4 that it is also possible to incorporate known random embedding results into our theoretical framework and thereby to obtain the best possible error decay rates that can be expected for $(P_{K,y})$ with quantized measurements. This achievement is due to a more general guarantee, Theorem 4.2, which is of interest in its own right.

The rest of this article is organized as follows: Section 2 presents our findings in full generality, where the proof of Theorem 2.7 is postponed to Section 6. Section 3 and 4 are then devoted to applications and specific examples, including a series of new guarantees for quantized compressed sensing; the corresponding proofs can be found in Section 7 and 8, respectively. Finally, some concluding remarks are made in Section 5.

1.3 Notation

The letters c and C are reserved for (positive) constants, whose values could change from time to time, and we speak of a *universal constant* if their values do 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 $A \leq C \cdot B$, and 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 $S \subset [p]$ is denoted by $|S|$ and its *set complement* in $[p]$ is given by $S^c := [p] \setminus S$. 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 B_q^p . The *Euclidean unit sphere* is given by $S^{p-1} := \{\mathbf{v} \in \mathbb{R}^p : \|\mathbf{v}\|_2 = 1\}$.

Let $H, H' \subset \mathbb{R}^p$ and $\mathbf{v} \in \mathbb{R}^p$. We write $\text{span}(H)$ and $\text{cone}(H)$ for the *linear hull* and *conic hull* of H , respectively, and $\chi_H : \mathbb{R}^p \rightarrow \{0, 1\}$ denotes its *indicator function*. 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 $P_v := P_{\text{span}(\{\mathbf{v}\})}$. 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 $H - \mathbf{v} := H - \{\mathbf{v}\}$. Moreover, 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 $\|\mathbf{a}\|_{\psi_2} := \sup_{\mathbf{v} \in S^{p-1}} \|\langle \mathbf{a}, \mathbf{v} \rangle\|_{\psi_2} < \infty$. Moreover, \mathbf{a} is *centered* if $\mathbb{E}[\mathbf{a}] = \mathbf{0}$, and it is *isotropic* if $\mathbb{E}[\mathbf{a}\mathbf{a}^\top] = \mathbf{I}_p$. Finally, we write $\mathbf{a} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ 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 [Ver18, 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

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

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 $\text{sign}(0) := 1$, and if $\text{sign}(\cdot)$ is applied to a vector, this operation is understood entrywise. Finally, the *ceiling* and *floor function* of $v \in \mathbb{R}$ are denoted by $\lceil v \rceil$ and $\lfloor v \rfloor$, respectively.

2 Main Result

As preliminary steps, we first introduce our formal model setup in Subsection 2.1, followed by several important increment conditions on the observation variable in Subsection 2.2. The main result of Theorem 2.7 is then stated and discussed in Subsection 2.3, while its proof can be found in Section 6.

2.1 Model Setup

The following model assumption fixes the notation for the remainder of this section and specifies the instance of Problem 1.1 that we will study from now on:

Assumption 2.1 (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 $L > 0$.³

³Note that we actually have that $L \geq \sqrt{1/\log(2)} > 1$ if \mathbf{a} is isotropic, cf. [JLPY20].

- *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) .
- (*Noiseless*) *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 problem:

- *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.

Let us briefly discuss the three parts of Assumption 2.1. Part (a) 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. The term ‘noiseless’ in part (b) does not imply that the $\tilde{y}_i(\mathbf{x})$ are deterministic when conditioned on \mathbf{a}_i . Indeed, F_i may be random but it should not be viewed as noise, since the measurement ensemble $\{(\mathbf{a}_i, F_i)\}_{i=1}^m$ is only drawn once and does not depend on a specific $\mathbf{x} \in \mathcal{X}$. However, we note that Theorem 2.7 allows for (adversarial) noise in the input of $(P_{K,y})$, in the sense that the observation vector $\mathbf{y} = (y_1, \dots, y_m)$ need not be equal to $\tilde{\mathbf{y}}(\mathbf{x}) = (\tilde{y}_1(\mathbf{x}), \dots, \tilde{y}_m(\mathbf{x}))$.⁴

Assumption 2.1(c) is probably the most abstract part of our model setup, but it is actually a crucial ingredient of our main result: Theorem 2.7 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). Indeed, the basic linearization strategy behind $(P_{K,y})$ typically calls for an appropriate transformation of the signal, due to non-linear observations; a good example is the special case Theorem 1.2, where T amounts to a rescaling of \mathbf{x} by the model-dependent factor μ . Such a simple “correction” will be sufficient for most applications considered in Section 3, but there exist more complicated situations, such as the variable selection problem, where \mathcal{X} contains support sets rather than vectors (see Subsection 3.5). 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 $(P_{K,y})$ and its constraint set $K \subset \mathbb{R}^p$.⁵

In principle, an arbitrary target function T is allowed in Assumption 2.1, but there often exists a canonical choice of T that is driven by the size of the following model parameter:

Definition 2.2 Let Assumption 2.1 be satisfied. We define the *target mismatch* of $\mathbf{x} \in \mathcal{X}$ as

$$\rho(\mathbf{x}) := \|P_{K_0}(\mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] - T\mathbf{x})\|_2,$$

where P_{K_0} is the Euclidean projection onto $K_0 := \text{cone}(K - T\mathcal{X})$.

As its name suggests, the target mismatch $\rho(\mathbf{x})$ measures the mismatch between the target vector $T\mathbf{x}$ and $\mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}]$, which is nothing else than the *global expected loss minimizer*, i.e., the solution to $\min_{\mathbf{z} \in \mathbb{R}^p} \mathbb{E}[(\tilde{\mathbf{y}}(\mathbf{x}) - \langle \mathbf{a}, \mathbf{z} \rangle)^2]$. In the context of our recovery guarantee Theorem 2.7, this expression can be seen as an upper bound for the asymptotic error of estimating $T\mathbf{x}$ via $(P_{K,y})$. Consequently, a general rule of thumb is to select $T\mathbf{x}$ such that $\rho(\mathbf{x})$ vanishes or is sufficiently small. However, simply setting $T\mathbf{x} := \mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}]$ (if contained in K) is not always consistent with our wish for signal recovery,

⁴The use of the modifier ‘ \sim ’ in Assumption 2.1 is also due to this fact and may help to distinguish between our probabilistic observation model and the actual input of $(P_{K,y})$.

⁵One should bear in mind that the target function T is a *theoretical* object that does not have any impact on the solution of $(P_{K,y})$ and may be (partially) unknown in practice. However, it is an essential component of our approach to Problem 1.1, relating the underlying observation model to the actual recovery method $(P_{K,y})$, see also Remark 2.8(1).




Figure 1: Illustration of the localization argument behind Fact 2.4. We assume that $\mathcal{E}_x(v) > 0$ for all $v \in K_{x,t}$, or equivalently, $\mathcal{E}_x(z - Tx) > 0$ for all $z \in K_{x,t} + Tx$ (see red arc). Since $\mathcal{E}_x(0) = 0$, the convexity of K and $\mathcal{E}_x(\cdot)$ implies that a minimizer \hat{z} of $(P_{K,y})$ must not lie outside the dark gray intersection, which simply means that $\|\hat{z} - Tx\|_2 \leq t$.

e.g., see Subsection 3.2. Therefore, we have left T unspecified in Assumption 2.1, even though a proper choice will be obvious for all examples of interest. The problem of choosing T is also closely related to the so-called *mismatch principle* developed in [Gen19] (see also the technical report [GK18]). We refer the reader to [Gen19, Chap. 3 & 4] for further reading.

Remark 2.3 The target mismatch $\rho(x)$ is clearly bounded by $\|\mathbb{E}[\tilde{y}(x)a] - Tx\|_2$, which is an even simpler expression. However, the projection P_{K_0} becomes useful in some situations, for instance, when Tx selects the *expected loss minimizer on K* , i.e., a solution to $\min_{z \in K} \mathbb{E}[(\tilde{y}(x) - \langle a, z \rangle)^2]$. In this case, we have that $\rho(x) = 0$ regardless of the observation model, which in turn ensures consistent estimation of Tx . This is in accordance with the fact that the empirical loss minimizer (i.e., a solution to $(P_{K,y})$) converges to the expected loss minimizer under mild conditions as $m \rightarrow \infty$; see [GK20, Def. 2.8 and Rmk. 2.9] for more details and possible refinements. \diamond

2.2 Analysis Approach and Increment Conditions

From now on, let Assumption 2.1 be satisfied. In this subsection, we introduce the key condition for our uniform recovery guarantee, namely that the class of observation variables $\{\tilde{y}(x)\}_{x \in \mathcal{X}}$, or at least an approximation thereof, has sub-Gaussian increments. To better understand the relevance of this condition, it is useful to first take a closer look at our basic analysis approach for the generalized Lasso:

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

Note that we hide the dependency on the measurement vectors $\{a_i\}_{i=1}^m$ when referring to $(P_{K,y})$, as this will be always clear from the context; in contrast, the input vector $y = (y_1, \dots, y_m) \in \mathbb{R}^m$ depends on the specific signal $x \in \mathcal{X}$ that one intends to recover via $(P_{K,y})$, i.e., y equals $\tilde{y}(x)$ or a noisy version thereof.

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

$$\mathcal{E}_x(v) := \tilde{\mathcal{L}}_x(v) - \tilde{\mathcal{L}}_x(0),$$

measuring how much the empirical loss is changing when traveling from Tx in direction of v . The following fact is an immediate consequence of the convexity of K and $\tilde{\mathcal{L}}_x$; see Figure 1 for an illustration.

Fact 2.4 For $x \in \mathcal{X}$ and a desired reconstruction accuracy $t > 0$, we set

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

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

For a fixed level of accuracy $t > 0$, Fact 2.4 implies *uniform* recovery for all those $x \in \mathcal{X}$ satisfying $\inf_{v \in K_{x,t}} \mathcal{E}_x(v) > 0$. This observation motivates us to show a uniform lower bound for the excess loss for all $x \in \mathcal{X}$ and $v \in K_{x,t}$. In particular, this task turns out to be considerably more difficult than showing a bound on $K_{x,t}$ for a fixed $x \in \mathcal{X}$, which would lead to a non-uniform error bound in accordance with the second branch of research discussed in Subsection 1.1. To get a better sense of this challenge, let us consider the following decomposition of the excess loss:

$$\mathcal{E}_x(v) = \underbrace{\frac{1}{m} \sum_{i=1}^m \langle a_i, v \rangle^2}_{=: \mathcal{Q}(v)} - \underbrace{\frac{2}{m} \sum_{i=1}^m (y_i - \tilde{y}_i(x)) \langle a_i, v \rangle}_{=: \mathcal{N}_x(v)} + \underbrace{\frac{2}{m} \sum_{i=1}^m (\langle a_i, Tx \rangle - \tilde{y}_i(x)) \langle a_i, v \rangle}_{=: \mathcal{M}_x(v)}. \quad (2.1)$$

The *quadratic term* $\mathcal{Q}(v)$ and *noise term* $\mathcal{N}_x(v)$ are rather unproblematic and can be controlled by a recent matrix deviation inequality (see Step 1 and Step 2 in Section 6). Much more intricate is the *multiplier term* $\mathcal{M}_x(v)$, since the underlying multiplier variable $\xi(x) := \langle a, Tx \rangle - \tilde{y}(x)$ depends on x , so that we actually have to deal with an *empirical product process*. In fact, a major difficulty is that known concentration results for such product processes do not apply directly, since the class $\{\xi(x)\}_{x \in \mathcal{X}}$ 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 1.2. The key idea of our approach is to approximate $\tilde{y}(x)$ by a more “regular” observation variable $\tilde{y}_t(x)$, depending on the desired accuracy t , in such a way that the resulting multiplier and the approximation error both have sub-Gaussian increments. This is made precise by the following assumption:

Assumption 2.5 Let $t \geq 0$ and define a pseudo-metric on \mathcal{X} by $d_T(x, x') := \|Tx - Tx'\|_2$ for $x, x' \in \mathcal{X}$. We assume that there exists a class of observation variables $\{\tilde{y}_t(x)\}_{x \in \mathcal{X}}$, such that the following properties hold:

(a) *Approximation error*: Setting $\varepsilon_t(x) := |\tilde{y}(x) - \tilde{y}_t(x)|$, we assume that $\mathbb{E}[\varepsilon_t(x) \cdot \langle a, z \rangle] \leq \frac{t}{64}$ for all $x \in \mathcal{X}$ and $z \in \mathbb{S}^{p-1}$.

(b) *Multiplier increments*: Set $\xi_t(x) := \langle a, Tx \rangle - \tilde{y}_t(x)$ for $x \in \mathcal{X}$. We assume that there exist $r \geq 0$ and $L_t \geq 0$ such that

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

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

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

Assumption 2.5(b) and (c) imply that $\{\xi_t(x)\}_{x \in \mathcal{X}}$ and $\{\varepsilon_t(x)\}_{x \in \mathcal{X}}$ have sub-Gaussian increments with respect to $L_t \cdot d_T$ and $\hat{L}_t \cdot d_T$, respectively, while r and \hat{r} bound the (sub-Gaussian) diameter of the respective class. A convenient interpretation of Assumption 2.5 is as follows: choose an approximation $\tilde{y}_t(x)$ for every $x \in \mathcal{X}$ such that the error does not get 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 parameters L_t and \hat{L}_t do not grow too rapidly as the desired accuracy t becomes smaller. A remarkable conclusion of our applications to quantized compressed sensing in Subsection 3.1–3.3 is that this strategy can even succeed for observations models with discontinuous output functions.

Finally, we would like to emphasize that the approximation error $\varepsilon_t(x)$ includes taking the absolute value, which is absolutely crucial to our proof (see Step 3 in Section 6). In particular, Assumption 2.5

does not necessarily imply that the original class $\{\zeta(x)\}_{x \in \mathcal{X}}$ 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 $\{\zeta(x)\}_{x \in \mathcal{X}}$ already has sub-Gaussian increments, such as in Theorem 1.2. In this case, we can simply choose $\tilde{y}_t(x) := \tilde{y}(x)$ so that $\varepsilon_t(x) = 0$ and Assumption 2.5(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 problems, e.g., see [ALMT14; CRPW12; MPT07; RV08; Sto09] for pioneering works in that direction. Our approach is no exception and makes use of the following localized version of mean width.

Definition 2.6 Let $H \subset \mathbb{R}^p$ and let $g \sim N(0, I_p)$. The (Gaussian) mean width of H is given by

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

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 will allow us to capture the size of measurement noise: For $m_0 \in \{0, 1, \dots, m\}$ and $w = (w_1, \dots, w_m) \in \mathbb{R}^m$, let

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

where (w_1^*, \dots, w_m^*) is the non-increasing rearrangement of $(|w_1|, \dots, |w_m|)$. Obviously, we have that $\|w\|_{[m_0]}^2 + \sigma_{m_0}(w)_2^2 = \|w\|_2^2$, and in particular, $\|w\|_{[0]} = 0$ and $\sigma_0(w)_2 = \|w\|_2$. 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 Section 3; see Section 6 for a complete proof.

Theorem 2.7 *There exist universal constants $c, C > 0$ for which the following holds.*

Let Assumption 2.1 and 2.5 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 that⁶

$$\begin{aligned} m \geq C \cdot L^2 \cdot \max \Big\{ & \log L \cdot (w_t^2(K - T\mathcal{X}) + u^2), \\ & L^2 \Delta^2 \cdot (w_t^2(K - T\mathcal{X}) + u_0^2), \\ & t^{-2}(r^2 + \hat{r}^2) \cdot (w_t^2(K - T\mathcal{X}) + u^2) + t^{-2}(L_t^2 + \hat{L}_t^2) \cdot w^2(T\mathcal{X}) \Big\}. \end{aligned} \quad (2.2)$$

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 $x \in \mathcal{X}$ with $\rho(x) \leq \frac{t}{32}$: Let $y \in \mathbb{R}^m$ be any input vector such that

$$\frac{1}{\sqrt{m}} \|y - \tilde{y}(x)\|_{[m_0]} \leq \Delta t \quad \text{and} \quad \frac{1}{\sqrt{m}} \sigma_{m_0}(y - \tilde{y}(x))_2 \leq \frac{t}{20}. \quad (2.3)$$

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

⁶In the case of exact recovery, i.e., $t = 0$, we follow the convention $0 \cdot \infty := 0$, so that the condition (2.2) requires that $r = \hat{r} = L_t = \hat{L}_t = 0$. Then, (2.2) is already 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 (2.3) imply that $(P_{K,y})$ is a *robust estimator* that allows for (adversarial) perturbations and outliers in the observation vector $\tilde{y}(x)$. For $m_0 = 0$, we obtain the simple baseline constraint $\frac{1}{\sqrt{m}}\|y - \tilde{y}(x)\|_2 \leq \frac{t}{20}$, which is consistent with standard noise bounds from the literature; note that in this case, the first branch of (2.2) simply absorbs the second one if $\Delta = L^{-1}\sqrt{\log L}$ and $u = u_0$. On the other hand, if $m_0 \geq 1$, the m_0 largest entries of the noise vector $y - \tilde{y}(x)$ in magnitude are treated as “gross” outliers. Indeed, choosing the free parameter Δ sufficiently large, the first condition of (2.3) becomes significantly less restrictive than the second one (concerning the remaining entries of $y - \tilde{y}(x)$); but at the same time, this relaxation may increase the number of required observations. The benefit of permitting outliers will become especially clear in the context of quantized measurement schemes where the noise is due to wrong bits in y instead of real-valued perturbations (see Subsection 3.1–3.3).

The second important constraint of Theorem 2.7 is that the error bound 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 of estimating Tx via $(P_{K,y})$, thereby preventing recovery up to arbitrarily high precision (cf. Remark 2.3).

Let us now turn to the key assumption (2.2), which relates the required number of measurements m to the desired accuracy t . The above formulation views m as a function of t , specifying how rapidly m grows if t decreases. The inverse relationship is also of considerable interest, since it describes the recovery error as a function of m . However, the latter can only be made explicit in specific situations where the parameters at the right-hand side of (2.2) 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 2.5, whose size can have a significant impact on the required (over-)sampling rate. For 1-bit observations as studied in Subsection 3.1, for example, we are able to achieve $L_t, \hat{L}_t \lesssim t^{-1}$, so that (2.2) 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 2.7 in Section 4 that can bypass the restrictions of Assumption 2.5 at the price of a stronger (local) stability condition on the observation model. This particularly allows us to combine our approach with available signal embedding results from the literature and thereby to obtain significantly stronger error decay rates in those cases.

The dependence of (2.2) on the complexity parameters $w_t(K - T\mathcal{X})$ and $w(T\mathcal{X})$ is quite common and natural. In fact, similar expressions have already appeared in several of the articles discussed in Subsection 1.1, e.g., see [DM18a; PV13b; PVY16]; we also refer to these works for examples of (upper) bounds in special cases, including the prototypical scenario of sparse recovery with an ℓ^1 -constraint. A detailed discussion of the mean width and possible simplifications is certainly an important topic in its own right, but would go beyond the scope of this paper. Nevertheless, a noteworthy feature of Theorem 2.7 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 more common parameter $w_t(K - K)$ can lead to improved guarantees in certain scenarios, especially when $T\mathcal{X}$ is much “smaller” than K ; see the part (2) of the following remark.

Remark 2.8 (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 2.7, one needs to be able to extract the actual signal $x \in \mathcal{X}$ from an approximation \hat{z} of Tx . At the present level of abstraction, we leave this additional step unspecified, and in fact, its implementation can differ considerably from situation to situation. In the case of single-index models (see Theorem 1.2), for example, this would involve a rescaling of \hat{z} , which would also require a (rough) knowledge of the non-linearity f . In contrast, for variable selection (see Subsection 3.5), it can be already sufficient to perform a simple hard-thresholding step on \hat{z} to obtain a good estimate of the underlying support.

(2) *Uniform vs. non-uniform recovery.* Except from the inclusion $T\mathcal{X} \subset K$, Theorem 2.7 does not impose any restrictions on the signal set \mathcal{X} . For example, if T is the identity function, the choice of $\mathcal{X} \subset K$ may range from a singleton $\mathcal{X} = \{x\}$ 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 $x \in K$, which is nothing else than non-uniform recovery. In particular, observing that $w(T\{x\}) = 0$, Theorem 2.7 is perfectly consistent with known non-uniform guarantees from the literature, e.g., see [Gen19; GK18]. Our main result therefore also indicates what additional expenses may come along with uniform recovery: On the one hand, we have to gain control over the increment parameters L_t and \hat{L}_t , which can lead to a weaker oversampling rate for non-linear observations; see Subsection 3.1–3.3 for concrete examples and Section 4 for a possible remedy. On the other hand, and more fundamentally, the behavior of the local mean width $w_t(K - T\mathcal{X})$ strongly depends on the size and shape of the signal set \mathcal{X} . Indeed, while $w_t(K - T\mathcal{X})$ is highly signal-dependent if $\mathcal{X} = \{x\}$, it rather captures the “worst-case” complexity over a whole signal class \mathcal{X} in the uniform case. Remarkably, even in relatively simple situations, either perspective can lead to very different complexity bounds, e.g., see [GMS20] for a recent result on total variation with gradient-sparse signals as well as [GKM20] for a study of the analysis formulation in compressed sensing. We believe that a deeper understanding of this observation is an interesting direction of further research, especially concerning intermediate cases where $T\mathcal{X}$ is an appropriate subset of K .

(3) *Optimality.* The best possible error decay rate that can result from Theorem 2.7 is $O(m^{-1/2})$, supposed that $L_t, \hat{L}_t \lesssim 1$.⁷ In fact, this rate cannot be improved *in general* (e.g., see [PVY16, Sec. 4]), or in other words, the exponent $-1/2$ is not an artifact of our proof but rather corresponds to a fundamental statistical barrier. However, it is possible to break through this barrier in *specific* model setups, such as *noiseless* 1-bit observations [JLBB13; JMPS19]. Having said that, such superior rates are usually not achieved by $(P_{K,y})$, but often require a more sophisticated estimator instead.

(4) *Possible extensions.* For the sake of clarity, we did not present the most general version of Theorem 2.7 that we could have shown. Nevertheless, there are several variations and generalizations that we expect to be practicable, such as different convex loss functions for $(P_{K,y})$, alternative formulations of the optimization problem, different error norms, anisotropic and heavier-tailed measurement vectors, or phase-retrieval-like problems. We refer to earlier works of the first author [Gen19; Gen17; GK20] for more technical details in this regard. One possible extension that is very specific to uniform recovery concerns the choice of pseudo-metric in Assumption 2.5. In principle, d_T could be replaced by an arbitrary pseudo-metric d on the signal set \mathcal{X} . The assertion and proof of Theorem 2.7 would remain literally 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 Section 6. While this relaxation would make our approach slightly more flexible, we note that it is difficult to control Talagrand’s γ_2 -functional efficiently in general. \diamond

3 Applications and Examples

This section demonstrates how to derive “out-of-the-box” guarantees from Theorem 2.7 for specific observation models. Subsection 3.1–3.3 are devoted to applications to quantized compressed sensing, while Subsection 3.4 revisits the case of single-index models. Subsection 3.5 is then concerned with a conceptually different example on the problem of variable selection. Note that all proofs for this section are deferred to Section 7, which contains more details on how Assumption 2.1 and 2.5 are precisely implemented in each case.

3.1 1-Bit Observations

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

$$y_i = \text{sign}(\langle a_i, x \rangle) + v_i, \quad i = 1, \dots, m. \quad (3.1)$$

⁷Note 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.




Figure 2: Illustration of the approximation strategy that is used in the proof of Corollary 3.1. 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 $\phi(s) = |\psi_t(s) - \text{sign}(s)|$ (plotted in dashed red) are then both $\frac{128}{t}$ -Lipschitz, which is already enough to fulfill Assumption 2.5 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.

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 (3.1) due to the scaling invariance of the sign-function, and the best one can hope for is to recover a normalized version of \mathbf{x} .

It turns out that the above 1-bit observation model is particularly compatible with *Gaussian* measurement vectors—a related result on general 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 $\mathbf{x} \in \mathcal{X}$, which implies that uniform recovery is possible up to arbitrarily high precision. The following guarantee makes this claim precise and is an application of Theorem 2.7 to Gaussian 1-bit observations. Its proof in Subsection 7.1 will demonstrate the usefulness of the approximation condition from Assumption 2.5, see also Figure 2 for an illustration of the underlying argument.

Corollary 3.1 *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 $\mathbf{a} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$. 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 that

$$m \geq C' \cdot t^{-2} \cdot \left(w_t^2(K - T\mathcal{X}) + t^{-2} \cdot w^2(T\mathcal{X}) + u^2 \right). \quad (3.2)$$

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 (3.1) such that $\frac{1}{2m} \|\mathbf{v}\|_1 \leq \beta$. Then every minimizer $\hat{\mathbf{z}}$ of $(\mathbf{P}_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \sqrt{\frac{2}{\pi}} \frac{\mathbf{x}}{\|\mathbf{x}\|_2}\|_2 \leq t$.

Corollary 3.1 is in line with some of the early achievements in 1-bit compressed sensing, for instance, see [PV13a; PV13b]. Remarkably, the condition (3.2) 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 [PV13b, Thm. 1.3]. In fact, we are not aware of any result in the literature that implies the statement of Corollary 3.1. But we stress that the above oversampling rate is still not optimal (cf. [JLBB13]) and can be further improved with a more specialized argument relying on random hyperplane tessellations (see Corollary 4.4 in Section 4). Finally, let us point out that the noise constraint of Corollary 3.1 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 the accuracy t (up to a log-factor). This condition again significantly improves [PV13b, Thm. 1.3] and is a particular consequence of the outlier robustness established in Theorem 2.7; see also [DM18a] for a similar achievement in the situation of dithered observations.

3.2 1-Bit Observations with Dithering

According to Ai et al. [ALPV14], the conclusion of Corollary 3.1 cannot be extended to sub-Gaussian measurements in general, regardless of the considered recovery method. A practicable remedy for this issue 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 [DK06; GN98; GS93], its benefits recently also emerged in compressed sensing theory [BFNPW17; DM18a; DM18b; KSW16; TR20; XJ18]. For more background information and recent advances, we refer to [Dir19] and the references therein.

Extending the original 1-bit model (3.1) by an additional dithering step leads to observations $\mathbf{y} \in \{-1, 1\}^m$ of the following form:

$$y_i = \text{sign}(\langle \mathbf{a}_i, \mathbf{x} \rangle + \tau_i) + v_i, \quad i = 1, \dots, m, \quad (3.3)$$

where $\mathbf{v} := (v_1, \dots, v_m) \in \{-2, 0, 2\}^m$ again models noise and the *dithering variables* τ_i are independent copies of a random variable τ that is uniformly distributed on the interval $[-\lambda, \lambda]$ for a certain $\lambda > 0$. Note that a major difference between dithering and additive noise is that the parameter λ is known and adjustable in practice (while the τ_i could be in principle unknown). Indeed, the following corollary of Theorem 2.7 is based on the fact that a careful choice of λ allows us to control the size of the (possibly non-vanishing) target mismatch $\rho(\mathbf{x})$, where $T\mathbf{x} := \lambda^{-1}\mathbf{x}$.

Corollary 3.2 *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 $\mathcal{X} \subset \text{RB}_2^p$ 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

$$\begin{aligned} \lambda &\geq \tilde{C} \cdot R \cdot L \cdot \sqrt{\log(e/t)}, \\ m &\geq C' \cdot L^2 \cdot \left((\log L + t^{-2}) \cdot (w_t^2(K - \lambda^{-1}\mathcal{X}) + u^2) + t^{-4} L^2 \lambda^{-2} \cdot w^2(\mathcal{X}) \right). \end{aligned} \quad (3.4)$$

Finally, let $\beta \in [0, 1]$ be such that $\beta \sqrt{\log(e/\beta)} \leq c_0 t L^{-2}$. 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 (3.3) such that $\frac{1}{2m} \|\mathbf{v}\|_1 \leq \beta$. Then every minimizer $\hat{\mathbf{z}}$ of $(P_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \lambda^{-1}\mathbf{x}\|_2 \leq t$.

Corollary 3.2 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 [TR20, Thm. IV.1]. In comparison to Corollary 3.1, the most notable improvement 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 3.1, we are not aware of any result in the literature that implies the statement of Corollary 3.2. Nevertheless, the oversampling rate of $O(m^{-1/4})$ promoted by (3.4)—ignoring an additional log-factor due to the choice of λ —can be further improved with a proof strategy that is specifically tailored to 1-bit observations with dithering, see Section 4.

3.3 Multi-Bit Observations

While 1-bit measurements are just an important extreme case of quantized compressed sensing, a considerable part of the literature deals with multi-bit observation models; once again, see [BJKS15;

Dir19] for a good introduction to this topic. 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 quantized observations $\mathbf{y} \in \delta\mathbb{Z}^m$ of the form

$$y_i = q_\delta(\langle \mathbf{a}_i, \mathbf{x} \rangle + \tau_i) + v_i, \quad i = 1, \dots, m, \quad (3.5)$$

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 $\delta > 0$, and 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 the variables v_i take values in $\delta\mathbb{Z}$. An application of Theorem 2.7 then yields the following uniform recovery guarantee for multi-bit observations with sub-Gaussian measurement vectors and dithering.

Corollary 3.3 *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

$$m \geq C' \cdot L^2 \cdot \left((\log L + t^{-2}\delta^2) \cdot (w_t^2(K - \mathcal{X}) + u^2) + L^2 t^{-4} \delta^2 \cdot w^2(\mathcal{X}) \right). \quad (3.6)$$

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 (3.5) such that at least one of the following conditions is fulfilled:

$$(a) \frac{1}{\sqrt{m}} \|\mathbf{v}\|_2 \leq \frac{t}{40} \quad \text{or} \quad (b) \frac{1}{\sqrt{m}} \|\mathbf{v}\|_2 \leq \frac{c_0 \sqrt{\delta t}}{L^2 \sqrt{\max\{1, \log(\delta e/t)\}}} \quad \text{and} \quad \frac{1}{m} \|\mathbf{v}\|_0 \leq \frac{t}{\delta}. \quad (3.7)$$

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

Corollary 3.3 can again be considered as a uniform version of a recent result on the generalized Lasso by Thrampoulidis and Rawat [TR20, Thm. III.1]. For a fixed quantizer resolution $\delta > 0$, the condition (3.6) translates into an error decay rate of $O(\sqrt{\delta} \cdot m^{-1/4})$. On the other hand, if we consider t as fixed and assume that $\delta \ll t$, then (3.6) 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 $(P_{K, \mathbf{y}})$ is the same as if the input vector would consist of linear measurements (cf. Corollary 3.4 below). Note that this behavior is perfectly consistent with the fact that the uniformly quantized observations (3.5) become linear as $\delta \rightarrow 0$. It is also worth pointing out that the constraint (b) in (3.7) implies a certain outlier robustness of $(P_{K, \mathbf{y}})$: if $t < \delta$ and only a fraction of t/δ bits are corrupted, then the normalized ℓ^2 -noise error may essentially scale in the order of $\sqrt{\delta t}$ instead of t .

To the best of our knowledge, Corollary 3.3 is a new result, but we stress that there exist uniform recovery guarantees for other programs than $(P_{K, \mathbf{y}})$ in the literature which imply better oversampling rates, e.g., see [JMPS19, Thm. 3] or [XJ18] in the case of structured signal sets. Nevertheless, we expect that Corollary 3.3 can be easily improved in that regard, using the strategy of Section 4 in conjunction with known uniform embedding results.

3.4 Single-Index Models and Linear Observations

The situation of single-index models was already considered in Theorem 1.2 in Subsection 1.2, as an appetizer for our more general approach. As pointed out there, this guarantee demonstrates that uniform recovery can be expected far beyond quantized measurement schemes, thereby upgrading an earlier non-uniform result by Plan and Vershynin [PV16, Thm. 1.9]. In view of the increment condition of Assumption 2.5, the Lipschitz continuity of f in Theorem 1.2 is certainly not necessary, but it indicates that a more “regular” observation variable can lead to a better error decay rate; note that the rate $O(m^{-1/2})$ resulting from (1.2) cannot be improved in general (see Remark 2.8(3)). Apart

from that, we wish to emphasize that (Lipschitz) continuous non-linearities are not only of academic interest but also appear in practical applications, for instance, as power amplifiers in sensor networks [GJ20].

In principle, Theorem 1.2 could be extended to sub-Gaussian measurement without any further effort, but in general, the target mismatch $\rho(\mathbf{x})$ would not vanish even if T is allowed to pick any scalar multiple of $\mathbf{x} \in \mathcal{X}$ (see [Gen19, Prop. 4.6]). However, it turns out that this issue does not arise in the prototypical case of linear observations:

Corollary 3.4 *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 $\|\mathbf{a}\|_{\psi_2} \leq L$. 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

$$m \geq C \cdot L^2 \cdot \log L \cdot (w_t^2(K - \mathcal{X}) + u^2). \quad (3.8)$$

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

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

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

Compared to all other results of this section, Corollary 3.4 does not involve any oversampling factors, and in particular, it turns into an exact recovery guarantee for $t = 0$. Although the statement of Corollary 3.4 bears resemblance with standard results from compressed sensing, there is an important difference concerning the required number of measurements: while the condition (3.8) promotes the local mean width $w_t(K - \mathcal{X})$ as the parameter to be analyzed, it is more common in compressed sensing theory to relate the sampling rate to deterministic conditions like the nullspace or restricted isometry property, e.g., see [FR13]. For the standard sparsity model in combination with an ℓ^1 -constraint, both approaches lead to very similar results, but their outcome can deviate drastically in more advanced situations, for instance, when using a (non-orthogonal) analysis transform [GKM20; GMS20], see also Remark 2.8(2).

3.5 Variable Selection

This subsection is concerned with an instance of Assumption 2.1 that is substantially different from the previous ones: For a fixed integer $s \leq p$ and output function $f: \mathbb{R}^p \rightarrow \mathbb{R}$, we ask for recovery of an index set $\mathcal{S} \subset [p]$ with $|\mathcal{S}| \leq s$ from observations $\mathbf{y} \in \mathbb{R}^m$ of the form

$$y_i = f(\mathbf{a}_{i,\mathcal{S}}) + v_i, \quad i = 1, \dots, m, \quad (3.9)$$

where $\mathbf{a}_{i,\mathcal{S}} \in \mathbb{R}^p$ is the coordinate projection of \mathbf{a}_i onto \mathcal{S} and $v_i \in \mathbb{R}$ models arbitrary noise. Most notably, the signals of interest are not parameters vectors in \mathbb{R}^p anymore but correspond to (small) index sets, specifying those entries 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 (unknown) set of active variables among all features variables in \mathbf{a}_i ; see [Gen19, Subsec. 4.2.4] for more details and a related model setup. 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 of active variables $\mathcal{S} \subset [p]$ with $|\mathcal{S}| \leq s$ from (non-linear) observations of the form (3.9)? Our next corollary of Theorem 2.7 provides a result in that direction under natural conditions on the output function f .

Corollary 3.5 *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

a 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

$$\frac{\alpha}{\sqrt{s}} \leq |(T\mathcal{S})_j| \leq \frac{\beta}{\sqrt{s}} \quad \text{for all } \mathcal{S} \in \mathcal{X} \text{ and } j \in \mathcal{S} \quad (3.10)$$

and

$$\|f(\mathbf{a}_{\mathcal{S}}) - f(\mathbf{a}_{\mathcal{S}'})\|_{\psi_2} \leq \gamma \sqrt{\frac{|\mathcal{S} \triangle \mathcal{S}'|}{s}} \quad \text{and} \quad \|f(\mathbf{a}_{\mathcal{S}})\|_{\psi_2} \leq \kappa \quad \text{for all } \mathcal{S}, \mathcal{S}' \in \mathcal{X}, \quad (3.11)$$

where $\mathcal{S} \triangle \mathcal{S}' \subset [p]$ 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

$$m \geq C \cdot L^2 \cdot \left((\log L + t^{-2}(L^2\beta^2 + \kappa^2)) \cdot (w_t^2(K - T\mathcal{X}) + u^2) + t^{-2}(L^2 + \gamma^2\alpha^{-2}) \cdot w^2(T\mathcal{X}) \right). \quad (3.12)$$

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 (3.9) 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 $(\mathbf{P}_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - T\mathcal{S}\|_2 \leq t$.

The assumptions (3.10) and (3.11) can be seen as natural balancing properties of the underlying observation model: (3.10) requires that the entries 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 (3.11) 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 3.5 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 and then use the corresponding indices as to estimate the true 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 in turn would lead to an undesirable factor of s in (3.12). It is certainly possible to show refined versions of Corollary 3.5 and we suspect that sharper error bounds can be obtained for variable selection by considering a different error measure than the ℓ^2 -norm (cf. Remark 2.8(4)). Nevertheless, a detailed elaboration would go beyond the scope of this article, and we confine ourselves with the above proof of concept.

4 Uniform Recovery Without Increment Conditions

We have seen in the course of Subsection 3.1–3.3 that the increment conditions of Assumption 2.5 can lead to inferior error decay rates in the situation of 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 an appropriate $\varepsilon > 0$ and to apply the *non-uniform* version of Theorem 2.7 to each $\mathbf{x} \in \mathcal{X}_\varepsilon$ separately (cf. Remark 2.8(2)). 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 (2.3) and essentially requires that close points in $T\mathcal{X}$ imply close observation vectors.

Assumption 4.1 Let Assumption 2.1 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 $1 - \eta$:

$$\sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \|T\mathbf{x} - T\mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{\sqrt{m}} \|\tilde{\mathbf{y}}(\mathbf{x}) - \tilde{\mathbf{y}}(\mathbf{x}')\|_{[2m_0]} \leq \frac{1}{2}\Delta t \quad \text{and} \quad \sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathcal{X} \\ \|T\mathbf{x} - T\mathbf{x}'\|_2 \leq \varepsilon}} \frac{1}{\sqrt{m}} \sigma_{m_0}(\tilde{\mathbf{y}}(\mathbf{x}) - \tilde{\mathbf{y}}(\mathbf{x}'))_2 \leq \frac{t}{40}. \quad (4.1)$$

The above strategy leads to the following general uniform recovery guarantee; see Section 8 for a detailed proof.

Theorem 4.2 *There exist universal constants $c, C, C_0 > 0$ for which the following holds.*

Let Assumption 2.1 and 4.1 be satisfied, let $r := \sup_{x \in \mathcal{X}} \|\langle \mathbf{a}, T\mathbf{x} \rangle - \tilde{\mathbf{y}}(x)\|_{\psi_2}$, and assume $\rho(x) \leq \frac{t}{32}$ for every $x \in \mathcal{X}$. For $u \geq 1$ and $u_0 \geq \sqrt{2m_0 \log(em/2m_0)}$, we assume that

$$m \geq C \cdot L^2 \cdot \left((\log L + t^{-2}r^2) \cdot \left(\sup_{x \in \mathcal{X}} w_t^2(K - T\mathbf{x}) + u^2 \right) + L^2 \Delta^2 \cdot \left(\sup_{x \in \mathcal{X}} w_t^2(K - T\mathbf{x}) + u_0^2 \right) \right) \quad (4.2)$$

and

$$\min\{u^2, u_0^2\} \geq C_0 \cdot \log \mathcal{N}(T\mathcal{X}, \varepsilon). \quad (4.3)$$

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

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

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

The statement of Theorem 4.2 strongly resembles the one of Theorem 2.7 with the important difference that the former does not rely on Assumption 2.5. In particular, the condition (4.2) does not depend on the increment parameters L_t and \hat{L}_t anymore, making it much less restrictive than (2.2). In that regard, it is worth noting that the global complexity of \mathcal{X} is now measured in terms of the covering number $\mathcal{N}(T\mathcal{X}, \varepsilon)$ in (4.3) instead of the mean width $w(T\mathcal{X})$. Apart from that, (4.2) establishes a refined local complexity measure, due to $\sup_{x \in \mathcal{X}} w_t^2(K - T\mathbf{x}) \leq w_t^2(K - T\mathcal{X})$. But nevertheless, the gain of Theorem 4.2 is obviously linked to the verification of Assumption 4.1, which is usually a highly non-trivial task.

We now demonstrate two applications of Theorem 4.2 in the prototypical case of 1-bit observations. In this specific situation, it turns out that Assumption 4.1 is very 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 [OR15] for noiseless, Gaussian 1-bit measurements (cf. Subsection 3.1). Hereafter, we agree on the shortcut notation $[H]_\varepsilon := (\frac{1}{\varepsilon}H - \frac{1}{\varepsilon}H) \cap B_2^p$ 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 B_2^p instead of \mathbb{S}^{p-1} .

Theorem 4.3 ([OR15, 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 4.3 with Theorem 4.2 yields an improved version of Corollary 3.1 (see Section 8 for a proof):

Corollary 4.4 *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 $\mathbf{a} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$. 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

$$m \geq C' \cdot \left(t^{-2} \cdot \left(\sup_{\mathbf{x} \in \mathcal{X}} w_t^2(K - T\mathbf{x}) + u^2 \right) + \varepsilon^2 t^{-3} \log^{3/2}(e/t) \cdot w^2([T\mathcal{X}]_\varepsilon) \right). \quad (4.4)$$

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 (3.1) such that $\frac{1}{2m} \|\mathbf{y}\|_1 \leq \frac{\beta}{2}$. Then every minimizer $\hat{\mathbf{z}}$ of $(P_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \sqrt{\frac{2}{\pi}} \frac{\mathbf{x}}{\|\mathbf{x}\|_2}\|_2 \leq 2t$.

In order 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 $\log \mathcal{N}(T\mathcal{X}, \varepsilon) \lesssim \varepsilon^{-2} \cdot w^2(T\mathcal{X})$, so that in the *worst case*, Corollary 4.4 would imply an error decay rate of $O(m^{-1/4})$ up to log factors. However, for many types of low-dimensional signal sets, such as subspaces or collections of 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 [OR15, Sec. 2].⁸ In this case, Corollary 4.4 would yield an error decay rate of $O(m^{-1/2})$ up to log factors, which is superior to the achievement of Corollary 3.1 and matches the best possible rate that can be expected for $(P_{K, \mathbf{y}})$ in general (cf. Remark 2.8(3)).

Our second application of Theorem 4.2 concerns the setup of sub-Gaussian 1-bit observations with dithering, as considered in Subsection 3.2. In this case, the stability condition of Assumption 4.1 can be related to a recent embedding result of Dirksen and Mendelson [DM18a], which is based on hyperplane tessellations for sub-Gaussian vectors:

Theorem 4.5 ([DM18a, Thm. 2.9]) *For every $L > 0$, there exist constants $c, \bar{c}, C, \bar{C} > 0$ only depending on L for which the following holds.*

Let $\mathcal{X} \subset RB_2^p$ for some $R > 0$ and let $A \in \mathbb{R}^{m \times p}$ 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 $\lambda \geq \bar{C} \cdot R$. In addition, suppose that 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 $1 - \exp(-cm\beta)$:

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

Combining Theorem 4.5 with Theorem 4.2 now yields an improved version of Corollary 3.2 (see Section 8 for a proof):

Corollary 4.6 *For every $L > 0$, there exist universal constants $c_0, C_0 > 0$ and constants $c, c', \bar{C}, C' > 0$ 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 $\mathcal{X} \subset RB_2^p$

⁸Here, it is important to note that, in contrast to the constraint set K , the (transformed) signal set $T\mathcal{X}$ may be highly non-convex.

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

$$\begin{aligned} \lambda &\geq \tilde{C} \cdot R \cdot \sqrt{\log(e/t)}, \\ m &\geq C' \cdot \left(t^{-2} \cdot \left(\sup_{\mathbf{x} \in \mathcal{X}} w_t^2(K - \lambda^{-1}\mathbf{x}) + u^2 \right) + \varepsilon^2 t^{-3} \log^{3/2}(e/t) \cdot w^2([(2\lambda)^{-1}\mathcal{X}]_\varepsilon) \right). \end{aligned} \quad (4.5)$$

Finally, let $\beta \in [0, 1]$ be such that $\beta \sqrt{\log(e/\beta)} \leq c_0 t / L^2$. 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 (3.3) such that $\frac{1}{2m} \|\mathbf{v}\|_1 \leq \frac{\beta}{2}$. Then every minimizer $\hat{\mathbf{z}}$ of $(P_{K, \mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \lambda^{-1}\mathbf{x}\|_2 \leq 2t$.

Analogously to the above situation of noiseless 1-bit observation, Corollary 4.6 implies an error 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 many structured signal sets. This reproduces a recent finding of Dirksen and Mendelson in [DM18a, Thm. 1.7], but for a different estimator.

We close this section by pointing out that the error bounds achieved by Corollary 4.4 and 4.6 do still not match the information-theoretic optimal rate of $O(m^{-1})$, e.g., see [DM18b; JLB13]. Nevertheless, we suspect that this gap is not an artifact of our proof, but rather due to a fundamental performance limit of the generalized Lasso $(P_{K, \mathbf{y}})$, which might be overcome with a more specialized recovery method, e.g., see [JMPS19].

5 Conclusion and Discussion

The purpose of this concluding section is to briefly review the results of this work in the light of our initial objectives as well as remaining challenges. We divide our discussion into several aspects that we consider particularly important in that regard:

- *Observation models.* Probably the greatest benefit of our approach to Problem 1.1 is its flexibility. The setup of Assumption 2.1 allows us to implement and analyze virtually every non-linear observation model that is conceivable for the program $(P_{K, \mathbf{y}})$, including the full spectrum of models mentioned in Subsection 1.1. In fact, the list of examples that we have seen in Section 3 is only a very small selection of possible applications. Importantly, each of these results is accompanied by a conceptually simple and systematic proof that does not require any deeper insight into the underlying model mechanisms; see also the proof template at the beginning of Section 7. However, while our methodology is a general path towards competitive benchmark guarantees, the resulting oversampling rates do not always match the best possible rates from the literature (due to the increment condition of Assumption 2.5). In Section 4, we have presented a workaround for this issue, based on an additional local stability assumption on the observation model. This strategy can indeed lead to near-optimal error bounds, but its gain certainly depends on the availability of a strong embedding result in every considered case.
- *Non-Gaussian measurements.* Our main result allows us to explore the impact of non-Gaussian measurements onto the performance of $(P_{K, \mathbf{y}})$, in particular, when consistent recovery can be expected and how an observation model could be modified in that respect. While relaxing the conditions on isotropy and sub-Gaussian tails in Assumption 2.1 is certainly practicable (cf. Remark 2.8(4)), our analysis is inherently limited to *independent* measurement vectors. This excludes more structured measurement ensembles, such as partial random circulant matrices or lossless expanders [FR13, Chap. 12 & 13]. On the other hand, such ensembles typically satisfy a variant of the restricted isometry property, which is the basis for some recent advances in quantized compressed sensing, e.g., see [DJR17; Fou17; JC17; XJ18]. These findings provide strong evidence that at least parts of our recovery results remain valid for a much larger class of measurement vectors. Thus,

any guarantee with structured measurements that meets the generality of Theorem 2.7 would be significant.

- *Complexity parameters.* The research on high-dimensional signal recovery problems over the last decade has shown that the (local) Gaussian mean width is a natural, yet accurate complexity measure for convex programs such as $(P_{K,y})$. A specific feature of our approach is the role played by the (not necessarily convex) signal set \mathcal{X} . As indicated in the context of Remark 2.8(2), one can view \mathcal{X} as a property of the underlying observation model that allows us to study any situation between non-uniform and (fully) uniform recovery; note that this refinement is also the key to the statement of Theorem 4.2. However, it is important to bear in mind that the Gaussian mean width is inherently difficult to estimate in any specific situation, except from relatively simple cases such as standard sparsity constraints or ℓ^1 -balls.
- *Non-convexity and data-driven priors.* Remarkably, the argument of Fact 2.4 is the only point in our proof that relies on the convexity of the constraint set K . In principle, it is possible to drop the convexity assumption on K by analyzing the projected gradient descent method as an algorithmic implementation of $(P_{K,y})$, e.g., see [ORS18; OS17; SO20]. However, the feasibility of this approach certainly depends on whether an efficient projection onto K is available.⁹ A modern line of research on signal processing with non-convex optimization advocates the use of data-driven priors, as a consequence of many recent advances on *generative models* in machine learning research; see [BJPD17; LS20] and the references therein for more details. Although the basic strategy behind these methods bears resemblance to ours, we believe that the complexity of learned signal priors leads to a very different mathematical problem, whose understanding is still in its infancy. More specifically, we suspect that common complexity measures like the Gaussian mean width or covering numbers are only useful for worst-case guarantees, whereas the analysis of non-generic scenarios seems to be out of reach with such tools. The latter point is particularly supported by the fact that even the situation of convex priors is not sufficiently well understood yet (cf. [GKM20; GMS20; MBKW20]).

6 Proof of the Main Result (Theorem 2.7)

Recall that according to Fact 2.4, it suffices to show that for all $x \in \mathcal{X}$ and $y \in \mathbb{R}^m$ satisfying $\rho(x) \leq \frac{t}{32}$ and (2.3), we have that

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

To this end, we make use of the decomposition in (2.1),

$$\mathcal{E}_x(v) = \mathcal{Q}(v) + \mathcal{N}_x(v) + \mathcal{M}_x(v),$$

and continue by showing separate bounds for every term, where each bound holds uniformly for all $x \in \mathcal{X}$ and $v \in K_{x,t}$. Unless stated otherwise, we assume that the hypotheses of Theorem 2.7 are fulfilled throughout this section and we have that $t > 0$ (the argumentation for $t = 0$ is in fact much simpler, see Step 4b). Moreover, we will use the notation $K_{\mathcal{X},t} := \cup_{x \in \mathcal{X}} K_{x,t} = (K - T\mathcal{X}) \cap tS^{p-1}$.

Step 1: Bounding the quadratic term. Let $A \in \mathbb{R}^{m \times p}$ denote the matrix with row vectors a_1, \dots, a_m . In order to control the random variable

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

⁹Note that 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 even NP-hard, e.g., see [ROV14].

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 6.1 ([JLPY20, Cor. 1.2]) *There exists a universal constant $C_Q > 0$ 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 $1 - 3 \exp(-u^2)$:

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

We now apply this result to $H := K_{\mathcal{X},t} \subset tS^{p-1}$ within the setting of Theorem 2.7: According to the first branch of the assumption (2.2), 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_Q), Theorem 6.1 implies that the following holds with probability at least $1 - 3 \exp(-u^2)$:

$$\frac{3t^2}{2} \geq \mathcal{Q}(\mathbf{v}) \geq \frac{t^2}{2} \quad \text{for all } \mathbf{v} \in K_{\mathcal{X},t}. \quad (6.1)$$

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

Theorem 6.2 ([DM18a, Thm. 2.10]) *There exist universal constants $c, C_N > 0$ 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 $1 - 2 \exp(-cu_0^2)$:

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

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

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

Let us now estimate the first summand of this bound. According to Theorem 6.2, there exist universal constants $c, C_N > 0$ such that for every $u_0 \geq \sqrt{m_0 \log(em/m_0)}$, we have the following with

probability at least $1 - 2 \exp(-cu_0^2)$:

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

where the last inequality follows from the second branch of the condition (2.2). Hence, taking a union bound with the event of (6.1), it follows that with probability at least $1 - 3 \exp(-u^2) - 2 \exp(-cu_0^2)$, the following holds uniformly for every $x \in \mathcal{X}$ and every $y \in \mathbb{R}^m$ satisfying the condition (2.3):

$$\sup_{v \in K_{x,t}} |\mathcal{N}_x(v)| \leq 2 \cdot \frac{t}{16} \cdot t + 2 \cdot \sqrt{\frac{3}{2}} \cdot t \cdot \frac{t}{20} \leq \frac{t^2}{4}. \quad (6.2)$$

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

$$\frac{1}{2} \mathcal{M}_x(v) = \frac{1}{m} \sum_{i=1}^m (\langle a_i, Tx \rangle - \tilde{y}_i(x)) \langle a_i, v \rangle$$

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

$$\frac{1}{2} \mathcal{M}_x(v) \geq \left(\inf_{v \in K_{x,t}} \frac{1}{2} \mathbb{E}[\mathcal{M}_x(v)] \right) - \frac{1}{2} |\mathcal{M}_x(v) - \mathbb{E}[\mathcal{M}_x(v)]|. \quad (6.3)$$

Since a is isotropic, we observe that $\frac{1}{2} \mathbb{E}[\mathcal{M}_x(v)] = \langle Tx - \mathbb{E}[\tilde{y}(x)a], v \rangle$. Furthermore, for every vector $z \in \mathbb{R}^p$ and every convex cone $H \subset \mathbb{R}^p$, the following estimate holds (see [ORS18, Lem. 16]):

$$\sup_{v \in H \cap S^{p-1}} \langle z, v \rangle \leq \|P_H z\|_2.$$

Combining these two facts with Definition 2.2 and the inclusion $\frac{1}{t} K_{x,t} \subset \text{cone}(K - Tx) \cap S^{p-1} = K_0 \cap S^{p-1}$, we obtain a lower bound in terms of the target mismatch:

$$\begin{aligned} \inf_{v \in K_{x,t}} \frac{1}{2} \mathbb{E}[\mathcal{M}_x(v)] &\geq t \cdot \inf_{v \in K_0 \cap S^{p-1}} \langle Tx - \mathbb{E}[\tilde{y}(x)a], v \rangle \\ &= -t \cdot \sup_{v \in K_0 \cap S^{p-1}} \langle \mathbb{E}[\tilde{y}(x)a] - Tx, v \rangle \geq -t \cdot \rho(x). \end{aligned} \quad (6.4)$$

We now turn to the centered multiplier term in (6.3). To this end, recall Assumption 2.5 and let $\{\tilde{y}_{t,i}(x)\}_{x \in \mathcal{X}}$ be independent copies of the class $\{\tilde{y}_t(x)\}_{x \in \mathcal{X}}$ for $i = 1, \dots, m$.¹⁰ Inserting these observations variables and applying the triangle inequality then yields

$$\begin{aligned} \frac{1}{2} |\mathcal{M}_x(v) - \mathbb{E}[\mathcal{M}_x(v)]| &\leq \frac{1}{m} \left| \sum_{i=1}^m \underbrace{(\langle a_i, Tx \rangle - \tilde{y}_{t,i}(x)) \langle a_i, v \rangle}_{=: \zeta_{t,i}(x)} - \mathbb{E}[(\langle a_i, Tx \rangle - \tilde{y}_{t,i}(x)) \langle a_i, v \rangle] \right| \\ &\quad + \mathbb{E}[\underbrace{|\tilde{y}(x) - \tilde{y}_t(x)|}_{=: \varepsilon_t(x)} \cdot |\langle a, v \rangle|] + \frac{1}{m} \sum_{i=1}^m \underbrace{|\tilde{y}_{t,i}(x) - \tilde{y}_t(x)|}_{=: \varepsilon_{t,i}(x)} \cdot |\langle a_i, v \rangle|. \end{aligned} \quad (6.5)$$

Note that the second empirical product process in (6.5) 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

¹⁰To be more precise, we assume that the tuples $(a_i, F_i, \tilde{y}_{t,i})$ are independent copies of (a, F, \tilde{y}_t) for $i = 1, \dots, m$. In particular, the respective conditions of Assumption 2.5(a)–(c) are also satisfied for $\{\tilde{y}_{t,i}(x)\}_{x \in \mathcal{X}}$.

and subtracting the expected value of the last summand in (6.5) and using Assumption 2.5(a), we obtain the following upper bound:

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

This estimate in conjunction with (6.4) implies that the following holds uniformly for all $x \in \mathcal{X}$ with $\rho(x) \leq \frac{t}{32}$:

$$\begin{aligned} \inf_{v \in K_{x,t}} \mathcal{M}_x(v) &\geq -\frac{t^2}{8} - 2 \sup_{v \in K_{x,t}} \frac{1}{m} \left| \sum_{i=1}^m \zeta_{t,i}(x) \langle a_i, v \rangle - \mathbb{E}[\zeta_{t,i}(x) \langle a_i, v \rangle] \right| \\ &\quad - 2 \sup_{v \in K_{x,t}} \frac{1}{m} \left| \sum_{i=1}^m \varepsilon_{t,i}(x) \cdot |\langle a_i, v \rangle| - \mathbb{E}[\varepsilon_{t,i}(x) \cdot |\langle a_i, v \rangle|] \right| \\ &\geq -\frac{t^2}{8} - 2 \sup_{\substack{x \in \mathcal{X} \\ v \in K_{x,t}}} \frac{1}{m} \left| \sum_{i=1}^m \zeta_{t,i}(x) \langle a_i, v \rangle - \mathbb{E}[\zeta_{t,i}(x) \langle a_i, v \rangle] \right| \\ &\quad - 2 \sup_{\substack{x \in \mathcal{X} \\ v \in K_{x,t}}} \frac{1}{m} \left| \sum_{i=1}^m \varepsilon_{t,i}(x) \cdot |\langle a_i, v \rangle| - \mathbb{E}[\varepsilon_{t,i}(x) \cdot |\langle a_i, v \rangle|] \right|. \end{aligned} \quad (6.6)$$

By Assumption 2.5(b) and (c), all factors of the product processes

$$\left\{ \frac{1}{m} \sum_{i=1}^m \zeta_{t,i}(x) \langle a_i, v \rangle \right\}_{x \in \mathcal{X}, v \in K_{x,t}}, \quad \left\{ \frac{1}{m} \sum_{i=1}^m \varepsilon_{t,i}(x) \cdot |\langle a_i, v \rangle| \right\}_{x \in \mathcal{X}, v \in K_{x,t}} \quad (6.7)$$

have sub-Gaussian increments. A key ingredient for controlling these empirical processes is a powerful concentration inequality due to Mendelson [Men16]; the following result is adapted from [Men16, Thm. 1.13].

Theorem 6.3 *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 $1 - 2\exp(-cu^2)$:

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

where $\gamma_2(\cdot, \cdot)$ denotes Talagrand's γ_2 -functional (e.g., see [Men16, Def. 1.2]).

Let us now apply Theorem 6.3 to the first product process in (6.7). 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 2.5(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 6.3 implies that the following holds with probability at least $1 - 2\exp(-cu^2)$:¹¹

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

According to Talagrand's Majorizing Measure Theorem [Tal14, Thm. 2.4.1], we have that $w(H) \asymp \gamma_2(H, \|\cdot\|_2)$ 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 $1 - 2\exp(-cu^2)$:

$$\begin{aligned} \sup_{\substack{\mathbf{x} \in \mathcal{X} \\ \mathbf{v} \in K_{\mathcal{X},t}}} \frac{1}{m} \left| \sum_{i=1}^m \xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle - \mathbb{E}[\xi_{t,i}(\mathbf{x}) \langle \mathbf{a}_i, \mathbf{v} \rangle] \right| \\ \leq C'_{\mathcal{M}} \cdot t \cdot L \cdot \left(\frac{(L_t w(T\mathcal{X}) + u \cdot r) \cdot (w_t(K - T\mathcal{X}) + u)}{m} \right. \\ \left. + \frac{r \cdot w_t(K - T\mathcal{X}) + L_t w(T\mathcal{X}) + u \cdot r}{\sqrt{m}} \right). \quad (6.8) \end{aligned}$$

The second product process in (6.7) 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 2.5(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 6.3 shows that with probability at least $1 - 2\exp(-cu^2)$, we have that

$$\begin{aligned} \sup_{\substack{\mathbf{x} \in \mathcal{X} \\ \mathbf{v} \in K_{\mathcal{X},t}}} \frac{1}{m} \left| \sum_{i=1}^m \varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle| - \mathbb{E}[\varepsilon_{t,i}(\mathbf{x}) \cdot |\langle \mathbf{a}_i, \mathbf{v} \rangle|] \right| \\ \leq C'_{\mathcal{M}} \cdot t \cdot L \cdot \left(\frac{(\hat{L}_t w(T\mathcal{X}) + u \cdot \hat{r}) \cdot (w_t(K - T\mathcal{X}) + u)}{m} \right. \\ \left. + \frac{\hat{r} \cdot w_t(K - T\mathcal{X}) + \hat{L}_t w(T\mathcal{X}) + u \cdot \hat{r}}{\sqrt{m}} \right). \quad (6.9) \end{aligned}$$

We are ready to prove our main result.

Step 4a: Proof of Theorem 2.7 ($t > 0$). We now assume that the events corresponding to (6.1), (6.2), (6.8), and (6.9) 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 (6.6), the following holds uniformly for every $\mathbf{x} \in \mathcal{X}$ with $\rho(\mathbf{x}) \leq \frac{t}{32}$

¹¹To be slightly more precise, we apply Theorem 6.3 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 $(\mathbf{a}, F, \tilde{y}_t)$.

and every $\mathbf{y} \in \mathbb{R}^m$ satisfying the condition (2.3):

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

If we could show that this lower bound is strictly positive, the claim of Theorem 2.7 would follow directly from Fact 2.4. To conclude this argument, it is enough to have that

$$m \geq C' \cdot L^2 t^{-2} \cdot \left((r + \hat{r}) \cdot w_t(K - T\mathcal{X}) + (L_t + \hat{L}_t) \cdot w(T\mathcal{X}) + u \cdot (r + \hat{r}) \right)^2, \quad (6.10)$$

$$m \geq C' \cdot L t^{-1} \cdot ((L_t + \hat{L}_t) \cdot w(T\mathcal{X}) + u \cdot (r + \hat{r})) \cdot (w_t(K - T\mathcal{X}) + u), \quad (6.11)$$

where $C' > 0$ is an appropriate universal constant. Indeed, both (6.10) and (6.11) are consequences of the condition (2.2): The bound of (6.10) is equivalent to the third branch of (2.2), while (6.11) follows from the multiplication of the first and third branch of (2.2) and then taking the square root.¹²

Step 4b: Proof of Theorem 2.7 ($t = 0$). In this case, 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 $\rho(\mathbf{x}) = 0$ (i.e., we have a noiseless linear model). Consequently, it follows that $\mathcal{N}_x(\cdot) = \mathcal{M}_x(\cdot) = 0$. Analogously to (6.1), 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 $\mathbf{v} \in K_{\mathcal{X},0} := \text{cone}(K - T\mathcal{X}) \cap S^{p-1}$. On this event, let $\hat{\mathbf{z}} \in K$ be any minimizer of $(\mathbf{P}_{K,\mathbf{y}})$ 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

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

which contradicts that fact that $\hat{\mathbf{z}}$ is a solution to $(\mathbf{P}_{K,\mathbf{y}})$. ■

7 Proofs for Section 3

Each of the results in Subsection 3.1–3.5 is an application of Theorem 2.7 and follows the same proof template:

Step 1. How the prerequisites of Theorem 2.7 are met:

Step 1a. Implementation of Assumption 2.1.

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

Step 1c. Controlling the increment parameters of Assumption 2.5.

Step 2. Proof of the actual statement via Theorem 2.7.

Step 3. Verification of Step 1b.

Step 4. Verification of Step 1c.

¹²Here we also use that 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} .

7.1 Proofs for Subsection 3.1

Proof of Corollary 3.1. We follow the proof template from the beginning of Section 7:

Step 1a. The model setup of Corollary 3.1 fits into Assumption 2.1 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 $T\mathbf{x} := \sqrt{\frac{2}{\pi}} \frac{\mathbf{x}}{\|\mathbf{x}\|_2}$. 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 $\tilde{y}_t(\mathbf{x})$ of the observation variable $\tilde{y}(\mathbf{x})$ such that the conditions of Assumption 2.5 are fulfilled with

$$L_t \lesssim 1 + t^{-1}, \quad \hat{L}_t \lesssim t^{-1}, \quad r \lesssim 1, \quad \hat{r} \lesssim 1. \quad (7.1)$$

Step 2. The first and third branch of the condition (2.2) are implied by (3.2) 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 (3.2) 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 (2.2) is satisfied if c_0 is chosen small enough and C' in (3.2) large enough.

Next, we show that every $\mathbf{y} \in \{-1, 1\}^m$ given by (3.1) with $\frac{1}{2m} \|\mathbf{v}\|_1 \leq \beta$ also satisfies the condition (2.3). Indeed, 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 $c_0 \leq 1/4$. Theorem 2.7 now implies the claim of Corollary 3.1.

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 $\langle \mathbf{a}, \bar{\mathbf{x}} \rangle \sim \mathcal{N}(0, 1)$ is independent of $P_{\mathbf{x}^\perp}(\mathbf{a})$, we have that

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

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

Step 4. Using the shortcut $\bar{x} := \frac{x}{\|x\|_2}$ again, we approximate the observation variable $\tilde{y}(x) = \text{sign}(\langle a, x \rangle)$ by $\tilde{y}_t(x) := \psi_t(\langle a, \bar{x} \rangle)$ for $t \in (0, 1]$ and $x \in \mathcal{X}$ 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 Figure 2 for an illustration. Since $\tilde{y}(x) = \text{sign}(\langle a, \bar{x} \rangle)$, the absolute value of the approximation error then takes the form

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

where we have used $t^\circ := t/128$ for the sake of notational convenience. We now show that for this choice of approximation, the conditions of Assumption 2.5 are indeed fulfilled with (7.1):

On Assumption 2.5(a). Let $x \in \mathcal{X}$, $z \in \mathbb{S}^{p-1}$, and consider the orthogonal decomposition

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

This implies

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

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

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

Moreover, by Jensen's inequality and the isotropy of a , it follows that

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

Putting everything together, this shows that Assumption 2.5(a) is satisfied.

On Assumption 2.5(b). Since ψ_t is $\frac{128}{t}$ -Lipschitz, the following holds for every $x, x' \in \mathcal{X}$ (with $\bar{x}' := \frac{x'}{\|x'\|_2}$):

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

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

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

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

On Assumption 2.5(c). We observe that the function $s \mapsto \phi(s) := |\text{sign}(s) - \psi_t(s)|$ is $\frac{128}{t}$ -Lipschitz. Therefore, for every $x, x' \in \mathcal{X}$, we obtain

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

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

7.2 Proofs for Subsection 3.2

Proof of Corollary 3.2. We follow the proof template from the beginning of Section 7:

Step 1a. The model setup of Corollary 3.2 fits into Assumption 2.1 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 $\mathcal{X} \subset RB_2^p$. 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

$$\lambda \geq \tilde{C} \cdot R \cdot L \cdot \sqrt{\log(e/t)}, \quad (7.2)$$

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 2.5 are fulfilled with

$$L_t \leq L \cdot (1 + 64t^{-1}), \quad \hat{L}_t \leq 64Lt^{-1}, \quad r \lesssim RL\lambda^{-1} + 1, \quad \hat{r} \lesssim 1. \quad (7.3)$$

Step 2. The first and third branch of the condition (2.2) are implied by (3.4) 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 $\beta\sqrt{\log(e/\beta)} = c_0 t L^{-2} \in (0, 1]$. 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 $u_0^2 \geq 2c_0 t m L^{-2}$. Combining the latter inequality with (3.4) for C' sufficiently large particularly implies that $u_0^2 \geq u^2$. Furthermore, we may assume that $c_0 \leq 1/4C$ and that (3.4) holds with $C' \geq 2C$, where $C > 0$ denotes the universal constant from (2.2). Then the second branch of the condition (2.2) 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

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

Next, we show that every $\mathbf{y} \in \{-1, 1\}^m$ given by (3.3) with $\frac{1}{2m}\|\mathbf{v}\|_1 \leq \beta$ also satisfies the condition (2.3). Indeed, 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 $c_0 \leq 1/4$. Theorem 2.7 now implies the claim of Corollary 3.2.

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

$$\rho(\mathbf{x}) = \|P_{K_0}(\mathbb{E}[\tilde{y}(\mathbf{x})\mathbf{a}] - \lambda^{-1}\mathbf{x})\|_2 \leq \|\mathbb{E}[(\text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \lambda^{-1}\mathbf{x} \rangle)\mathbf{a}]\|_2.$$

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:

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

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

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

where the last estimate is due to (7.2) 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 $t^\circ := t/64$ for the sake of notational convenience. We now show that for this choice of approximation, the conditions of Assumption 2.5 are indeed fulfilled with (7.3):

On Assumption 2.5(a). For $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathbb{S}^{p-1}$, we have that

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

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 2.5(a).

On Assumption 2.5(b). Since ψ_t is $\frac{64}{t\lambda}$ -Lipschitz, the following holds for every $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$:

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

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 2.5(c). We observe that the function $s \mapsto \phi(s) := |\text{sign}(s) - \psi_t(s)|$ 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 $\hat{r} \lesssim 1$. \blacksquare

7.3 Proofs for Subsection 3.3

Proof of Corollary 3.3. We follow the proof template from the beginning of Section 7:

Step 1a. The model setup of Corollary 3.3 fits into Assumption 2.1 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 2.5 are fulfilled with

$$L_t \lesssim L\delta t^{-1}, \quad \hat{L}_t \lesssim L\delta t^{-1}, \quad r \lesssim \delta, \quad \hat{r} \lesssim \delta. \quad (7.4)$$

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 $\tilde{y}_t(\mathbf{x}) := \tilde{y}(\mathbf{x})$, the conditions of Assumption 2.5 are fulfilled with respect to \mathcal{X}_t with

$$L_t \lesssim \delta t^{-1}, \quad \hat{L}_t = 0, \quad r \lesssim \delta, \quad \hat{r} = 0. \quad (7.5)$$

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 true provided that $c_0 \leq 1/40$. Furthermore, the condition $\frac{1}{m} \|\nu\|_0 \leq t\delta^{-1}$ is trivially fulfilled. Consequently, it suffices to prove the statement of Corollary 3.3 for all those input vectors $\mathbf{y} \in \delta\mathbb{Z}^m$ that satisfy the condition (a) in (3.7), and all those $\mathbf{y} \in \delta\mathbb{Z}^m$ that satisfy condition (b) in (3.7) 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 (7.4) it follows that the first and third branch of the condition (2.2) are implied by (3.6) 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 (2.2) follows from the first one. Therefore, the claim of Corollary 3.3 follows from Theorem 2.7 for all those input vectors $\mathbf{y} \in \delta\mathbb{Z}^m$ that satisfy the condition (a) in (3.7).

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

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

where $C > 0$ denotes the universal constant from (2.2). 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 (3.6) is large enough. Furthermore, the second branch of (2.2) is satisfied: we have that $CL^4 \Delta^2 u_0^2 = \frac{m}{2}$, and assuming that (3.6) holds with $C' \geq 1/2$, 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 (3.5) such that

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

also satisfies the condition (2.3). Indeed, since $\nu = \mathbf{y} - \tilde{\mathbf{y}}(\mathbf{x}) \in \delta\mathbb{Z}^m$ and $\|\nu\|_0 \leq tm\delta^{-1}$, it follows that $|\{i \in [m] : y_i \neq \tilde{y}_i(\mathbf{x})\}| \leq \lfloor tm\delta^{-1} \rfloor$. 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 2.7 now implies the claim of Corollary 3.3 for all those input vectors $\mathbf{y} \in \delta\mathbb{Z}^m$ that satisfy the condition (b) in (3.7).

The case $t > 128\delta$. Setting $m_0 := 0$, $\Delta := L^{-1} \sqrt{\log L}$, and $u := u_0$, the second branch of (2.2) is implied by the first one. Due to (7.5), Theorem 2.7 applied to the signal set \mathcal{X}_t now implies that there exist universal constants $c_1, C_1 > 0$ such that if

$$m \geq C_1 \cdot L^2 \cdot \left((\log L + t^{-2} \delta^2) \cdot (w_t^2(K - \mathcal{X}_t) + u^2) + t^{-4} \delta^2 \cdot w^2(\mathcal{X}_t) \right), \quad (7.6)$$

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 $\frac{1}{\sqrt{m}} \|\mathbf{y} - q_\delta(A\mathbf{x}' + \boldsymbol{\tau})\|_2 \leq \frac{t}{20}$, every minimizer $\hat{\mathbf{z}}$ of $(P_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - \mathbf{x}'\|_2 \leq t$.

Moreover, according to Theorem 6.1, there exists a universal constant $C_2 > 0$ such that if

$$m \geq C_2 \cdot L^2 \log L \cdot (w_t^2(\mathcal{X} - \mathcal{X}) + u^2), \quad (7.7)$$

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

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

We claim that Corollary 3.3 holds true 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 $\frac{1}{\sqrt{m}}\|\mathbf{y} - q_\delta(\mathbf{Ax} + \boldsymbol{\tau})\|_2 \leq \frac{t}{40}$. 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 using the triangle inequality, we obtain

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

On the event \mathcal{A}_1 , every minimizer $\hat{\mathbf{z}}$ of $(\mathbf{P}_{\mathbf{K}, \mathbf{y}})$ 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 (3.6) implies both (7.6) and (7.7). 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

$$\rho(\mathbf{x}) = \|P_{\mathbf{K}_0}(\mathbb{E}[\tilde{\mathbf{y}}(\mathbf{x})\mathbf{a}] - \mathbf{x})\|_2 \leq \|\mathbb{E}[(q_\delta(\langle \mathbf{a}, \mathbf{x} \rangle + \tau) - \langle \mathbf{a}, \mathbf{x} \rangle)\mathbf{a}]\|_2,$$

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

$$\mathbb{E}[q_\delta(s + \tau)] = s \quad \text{for all } s \in \mathbb{R}. \quad (7.8)$$

In other words, integrating over the dithering variable τ allows us to “eliminate” the quantizer. For the sake of completeness, let us show the identity (7.8) 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

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

Since τ and \mathbf{a} are independent, we can apply (7.8) 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 $t^\circ := t/128$ and approximate $\tilde{y}(x) = q_\delta(\langle a, x \rangle + \tau)$ by $\tilde{y}_t(x) := \psi_{t^\circ}(\langle a, x \rangle + \tau)$. The absolute value of the approximation error is then given by $\varepsilon_t(x) = \phi_{t^\circ}(\langle a, x \rangle + \tau)$. We now show that for this choice of approximation, the conditions of Assumption 2.5 are indeed fulfilled with (7.4):

On Assumption 2.5(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

$$\mathbb{E}_\tau[\chi_{\mathcal{E}_{t^\circ}}(s + \tau)] = \mathbb{E}_\tau[\chi_{\mathcal{E}_{t^\circ}}(\tau)] = \mathbb{E}_\tau[\chi_{\mathcal{E}_{0,t^\circ}}(\tau)] = \mathbb{E}_\tau[\chi_{[-t^\circ, t^\circ]}(\tau)] = t^\circ \delta^{-1}. \quad (7.9)$$

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 inequality (7.9), 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 2.5(a) is satisfied.

On Assumption 2.5(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 2.5(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 2.5 are fulfilled with respect to \mathcal{X}_t with (7.5). Since $\varepsilon_t(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathcal{X}_t$, Assumption 2.5(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:

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

This shows that Assumption 2.5(b) is satisfied with respect to \mathcal{X}_t with $L_t \lesssim \delta t^{-1}$ and $r \lesssim \delta$. ■

7.4 Proofs for Subsection 3.4

Proof of Theorem 1.2. We follow the proof template from the beginning of Section 7:

Step 1a. The model setup of Theorem 1.2 fits into Assumption 2.1 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 for $g \sim \mathcal{N}(0, 1)$, i.e., $Tx := \mu x$. In particular, we have that $d_T(x, x') = \mu \|x - x'\|_2$.

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

Step 1c. Let $g \sim \mathcal{N}(0, 1)$. For the trivial choice $\tilde{y}_t(x) := \tilde{y}(x)$, the conditions of Assumption 2.5 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 1.2 follows directly from Theorem 2.7.

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

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

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

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

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

Step 4. We simply set $\tilde{y}_t(x) := \tilde{y}(x)$. Then $\varepsilon_t(x) = 0$, implying that Assumption 2.5(a) and (c) are trivially fulfilled with $\hat{L}_t = \hat{r} = 0$. Furthermore, the following holds for every $x, x' \in \mathcal{X}$:

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

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

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

which shows that $r := \|f(g) - \mu g\|_{\psi_2}$ is a valid choice. Hence, Assumption 2.5(b) is satisfied as well. \blacksquare

Proof of Corollary 3.4. We follow the proof template from the beginning of Section 7:

Step 1a. The model setup of Corollary 3.4 fits into Assumption 2.1 as follows:

- The measurement vector $a \in \mathbb{R}^p$ is centered, isotropic, and sub-Gaussian with $\|a\|_{\psi_2} \leq L$. 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 $F(a, x) := \langle a, x \rangle$.
- The target function $T: \mathcal{X} \rightarrow K$ is the canonical embedding into K , i.e., $Tx := x$. In particular, we have that $d_T(x, x') = \|x - x'\|_2$.

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

Step 1c. For the trivial choice $\tilde{y}_t(x) := \tilde{y}(x)$, the conditions of Assumption 2.5 are fulfilled with

$$L_t = \hat{L}_t = r = \hat{r} = 0.$$

Step 2. Setting $m_0 := 0$, $\Delta := L^{-1} \sqrt{\log L}$, and $u := u_0$, the claim of Corollary 3.4 follows directly from Theorem 2.7.

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 $\tilde{\mathbf{y}}_t(\mathbf{x}) := \tilde{\mathbf{y}}(\mathbf{x})$. Then $\varepsilon_t(\mathbf{x}) = 0$, implying that Assumption 2.5(a) and (c) are trivially fulfilled with $\hat{L}_t = \hat{r} = 0$. Furthermore, since $\xi_t(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle - \tilde{\mathbf{y}}_t(\mathbf{x}) = 0$, Assumption 2.5(b) holds with $L_t = r = 0$. ■

7.5 Proofs for Subsection 3.5

Proof of Corollary 3.5. We follow the proof template from the beginning of Section 7:

Step 1a. The model setup of Corollary 3.5 fits into Assumption 2.1 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 [Ver18, 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 $\tilde{\mathbf{y}}_t(\mathcal{S}) := \tilde{\mathbf{y}}(\mathcal{S})$, the conditions of Assumption 2.5 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 3.5 follows directly from Theorem 2.7.

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

Step 4. We simply set $\tilde{\mathbf{y}}_t(\mathcal{S}) := \tilde{\mathbf{y}}(\mathcal{S})$. Then $\varepsilon_t(\mathcal{S}) = 0$, implying that Assumption 2.5(a) and (c) are trivially fulfilled with $\hat{L}_t = \hat{r} = 0$. To see that Assumption 2.5(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 (3.10), it follows that $\text{supp}(T\mathcal{S}) = \mathcal{S}$ for all $\mathcal{S} \in \mathcal{X}$. Using the lower bound in (3.10) again, we obtain the following estimate for all $\mathcal{S}, \mathcal{S}' \in \mathcal{X}$:

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

Combining this with the assumption (3.11), 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 (3.10) therefore yields $\|T\mathcal{S}\|_2 \leq \beta$ for every $\mathcal{S} \in \mathcal{X}$. Combining this estimate with (3.11), 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, $r \leq L\beta + \kappa$ is a valid choice for Assumption 2.5(b). ■

8 Proofs for Section 4

Proof of Theorem 4.2. 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 (4.2) and (4.3) allow us to first apply Theorem 2.7 to every $\mathbf{x}' \in \mathcal{X}_\varepsilon$ as a singleton signal set (Assumption 2.5 is trivially fulfilled here; cf. Remark 2.8(2)), 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

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

every minimizer $\hat{\mathbf{z}}$ of $(\mathbf{P}_{K,\mathbf{y}})$ satisfies $\|\hat{\mathbf{z}} - T\mathbf{x}'\|_2 \leq t$; note that we have applied the robustness conditions of (2.3) for $2m_0$ instead of m_0 here, which is possible due to $2m_0 \leq m$.

Now, let \mathcal{A}_2 denote the event of (4.1) in Assumption 4.1. 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 4.2 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 $\|\mathbf{x} - \mathbf{x}'\|_2 \leq \varepsilon$. 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 $(\mathbf{P}_{K,\mathbf{y}})$ 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 4.4. Analogously to Corollary 3.1, the model setup of Corollary 4.4 fits into Assumption 2.1 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 $F(\mathbf{a}, \mathbf{x}) := \text{sign}(\langle \mathbf{a}, \mathbf{x} \rangle)$. 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 the course of Subsection 7.1 (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

$$\beta \sqrt{\log(e/\beta)} = c_0 t \in (0, 1]. \quad (8.1)$$

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 (4.3) is fulfilled.

Similarly, it is not hard to see that the condition of (4.2) follows from (4.4) for C' sufficiently large (cf. Subsection 7.1 (Step 2)).

According to Theorem 4.3 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 4.3) such that if

$$m \geq C \cdot \left(\varepsilon^2 \beta^{-3} \cdot w^2([T\mathcal{X}]_\varepsilon) + \beta^{-1} \cdot \log \mathcal{N}(T\mathcal{X}, \varepsilon) \right) \quad (8.2)$$

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

$$\sup_{\substack{x, x' \in \mathcal{X} \\ \|Tx - Tx'\|_2 \leq \varepsilon}} \frac{1}{2m} \|\text{sign}(Ax) - \text{sign}(Ax')\|_1 \leq \frac{\beta}{2}.$$

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

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

and

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

where the last inequality holds for $c_0 > 0$ small enough. Consequently, Assumption 4.1 would hold for $\eta := \exp(-cm\beta)$ if we can show that (8.2) is satisfied under the hypotheses of Corollary 4.4. To this end, we first note that the relationship (8.1) 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 (4.4), 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, (8.2) is a consequence of (4.4) and the assumption $u^2 \geq C_0 \cdot \log \mathcal{N}(T\mathcal{X}, \varepsilon)$.

Since all assumptions of Theorem 4.2 are satisfied, the following holds with probability at least $1 - \exp(-cu^2)$ uniformly for every $x \in \mathcal{X}$: For any input vector $y \in \mathbb{R}^m$ such that

$$\frac{1}{\sqrt{m}} \|y - \tilde{y}(x)\|_{[2m_0]} \leq \frac{1}{2}\Delta t \quad \text{and} \quad \frac{1}{\sqrt{m}} \sigma_{m_0}(y - \tilde{y}(x))_2 \leq \frac{t}{40}, \quad (8.3)$$

every minimizer \hat{z} of $(P_{K,y})$ satisfies $\|\hat{z} - Tx\|_2 \leq t + \varepsilon \leq 2t$. The claim of Corollary 4.4 now follows from that fact that any input vector $y \in \{-1, 1\}^m$ given by (3.1) with $\frac{1}{2m} \|v\|_1 \leq \frac{\beta}{2}$ satisfies (8.3). ■

Proof of Corollary 4.6. Analogously to Corollary 3.2, the model setup of Corollary 4.6 fits into Assumption 2.1 as follows:

- The measurement vector $a \in \mathbb{R}^p$ is centered, isotropic, and sub-Gaussian with $\|a\|_{\psi_2} \leq L$. The signal set \mathcal{X} satisfies $\mathcal{X} \subset RB_2^p$. The output function $F: \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $F(a, x) = \text{sign}(\langle a, x \rangle + \tau)$, where τ is uniformly distributed on $[-\lambda, \lambda]$ and independent of a . In particular, F is a random function. Moreover, the observation vector of x is given by $\tilde{y}(x) = \text{sign}(Ax + \tau)$, where $A \in \mathbb{R}^{m \times p}$ denotes the sub-Gaussian random matrix with row vectors $a_1, \dots, a_m \in \mathbb{R}^p$ and $\tau := (\tau_1, \dots, \tau_m)$.

- 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 Subsection 7.2 (Step 3), there exists a universal constant $\tilde{C}' \geq e$ such that if

$$\lambda \geq \tilde{C}' \cdot R \cdot L \cdot \sqrt{\log(e/t)}, \quad (8.4)$$

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

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

for every $\mathbf{x} \in \mathcal{X}$, and together with (8.4), 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

$$\beta\sqrt{\log(e/\beta)} = c_0 t L^{-2} \in (0, 1]. \quad (8.5)$$

Now, we set

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

implying that $u_0 \geq \sqrt{2m_0\log(em/2m_0)}$, $m_0 \in \{0, 1, \dots, \lfloor \frac{m}{2} \rfloor\}$, and $u_0^2 \geq 2c_0 t m L^{-2}$. Combining the latter inequality with (4.5) 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 (4.3) is fulfilled. Similarly, it is not hard to see that the condition of (4.2) follows from (4.5) for $C' \gtrsim L^2 \log L$ sufficiently large (cf. Subsection 7.2 (Step 2)).

According to Theorem 4.5, there exist constants $c, \bar{c}, C, \bar{C} > 0$ only depending on L (possibly slightly different from those in Theorem 4.5) such that if $\lambda \geq \bar{C} \cdot R$ and

$$m \geq C \cdot \left(\varepsilon^2 \beta^{-3} \cdot w^2([(2\lambda)^{-1}\mathcal{X}]_\varepsilon) + \beta^{-1} \cdot \log \mathcal{N}(\lambda^{-1}\mathcal{X}, \varepsilon) \right) \quad (8.6)$$

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

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

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 4.1 would hold for $\eta := \exp(-cm\beta)$ if we can show that (8.6) is satisfied under the hypotheses of Corollary 4.6. To this end, we first note that the relationship (8.5) 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 (4.5), 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, (8.6) is a consequence of (4.5) and the assumption $u^2 \geq C_0 \cdot \log \mathcal{N}(\lambda^{-1}\mathcal{X}, \varepsilon)$.

Since all assumptions of Theorem 4.2 are satisfied, the following holds with probability at least $1 - \exp(-cu^2)$ uniformly for every $x \in \mathcal{X}$: For any input vector $y \in \mathbb{R}^m$ such that

$$\frac{1}{\sqrt{m}} \|y - \tilde{y}(x)\|_{[2m_0]} \leq \frac{1}{2} \Delta t \quad \text{and} \quad \frac{1}{\sqrt{m}} \sigma_{m_0}(y - \tilde{y}(x))_2 \leq \frac{t}{40}, \quad (8.7)$$

every minimizer \hat{z} of $(P_{K,y})$ satisfies $\|\hat{z} - \lambda^{-1}x\|_2 \leq t + \varepsilon \leq 2t$. The claim of Corollary 4.6 now follows from that fact that any input vector $y \in \{-1, 1\}^m$ given by (3.3) with $\frac{1}{2m} \|v\|_1 \leq \frac{\beta}{2}$ satisfies (8.7). ■

Acknowledgments

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

References

- [ALPV14] A. Ai, A. Lapanowski, Y. Plan, and R. Vershynin. “One-bit compressed sensing with non-Gaussian measurements”. *Linear Algebra Appl.* 441 (2014), 222–239.
- [ALMT14] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp. “Living on the edge: phase transitions in convex programs with random data”. *Inf. Inference* 3.3 (2014), 224–294.
- [BFNPW17] R. G. Baraniuk, S. Foucart, D. Needell, Y. Plan, and M. Wotter. “Exponential Decay of Reconstruction Error From Binary Measurements of Sparse Signals”. *IEEE Trans. Inf. Theory* 63.6 (2017), 3368–3385.
- [BJPD17] A. Bora, A. Jalal, E. Price, and A. G. Dimakis. “Compressed Sensing using Generative Models”. *Proceedings of the 34th International Conference on Machine Learning (ICML)*. Ed. by D. Precup and Y. W. Teh. Vol. 70. 2017, 537–546.
- [BJS15] P. T. Boufounos, L. Jacques, F. Krahmer, and R. Saab. “Quantization and Compressive Sensing”. *Compressed Sensing and its Applications: MATHEON Workshop 2013*. Ed. by H. Boche, R. Calderbank, G. Kutyniok, and J. Vybiral. Applied and Numerical Harmonic Analysis. Birkhäuser Cham, 2015, 193–237.
- [Bri82] D. R. Brillinger. “A Generalized Linear Model With “Gaussian” Regressor Variables”. *A Festschrift For Erich L. Lehmann*. Ed. by P. J. Bickel, K. Doksum, and J. Hodges. Chapman and Hall/CRC, 1982, 97–114.
- [CRT06a] E. J. Candès, J. K. Romberg, and T. Tao. “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information”. *IEEE Trans. Inf. Theory* 52.2 (2006), 489–509.
- [CRT06b] E. J. Candès, J. K. Romberg, and T. Tao. “Stable signal recovery from incomplete and inaccurate measurements”. *Comm. Pure Appl. Math.* 59.8 (2006), 1207–1223.
- [CRPW12] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky. “The Convex Geometry of Linear Inverse Problems”. *Found. Comput. Math.* 12.6 (2012), 805–849.
- [DK06] O. Dabeer and A. Karnik. “Signal Parameter Estimation Using 1-Bit Dithered Quantization”. *IEEE Trans. Inf. Theory* 52.12 (2006), 5389–5405.
- [Dir19] S. Dirksen. “Quantized Compressed Sensing: A Survey”. *Compressed Sensing and Its Applications: Third International MATHEON Conference 2017*. Ed. by H. Boche, G. Caire, R. Calderbank, G. Kutyniok, R. Mathar, and P. Petersen. Applied and Numerical Harmonic Analysis. Birkhäuser Cham, 2019, 67–95.
- [DJR17] S. Dirksen, H. C. Jung, and H. Rauhut. “One-bit compressed sensing with partial Gaussian circulant matrices”. Preprint arXiv:1710.03287. 2017.
- [DM18a] S. Dirksen and S. Mendelson. “Non-Gaussian Hyperplane Tessellations and Robust One-Bit Compressed Sensing”. Preprint arXiv:1805.09409. 2018.
- [DM18b] S. Dirksen and S. Mendelson. “Robust one-bit compressed sensing with partial circulant matrices”. Preprint arXiv:1812.06719. 2018.

- [Don06] D. L. Donoho. “Compressed sensing”. *IEEE Trans. Inf. Theory* 52.4 (2006), 1289–1306.
- [Fou17] S. Foucart. “Flavors of Compressive Sensing”. *Approximation Theory XV: San Antonio 2016*. Ed. by G. E. Fasshauer and L. L. Schumaker. Springer Proceedings in Mathematics & Statistics. 2017, 61–104.
- [FR13] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis. Birkhäuser Basel, 2013.
- [Gen19] M. Genzel. “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: <http://dx.doi.org/10.14279/depositonce-8394>. Technische Universität Berlin, 2019.
- [Gen17] M. Genzel. “High-Dimensional Estimation of Structured Signals From Non-Linear Observations With General Convex Loss Functions”. *IEEE Trans. Inf. Theory* 63.3 (2017), 1601–1619.
- [GJ20] M. Genzel and P. Jung. “Recovering Structured Data From Superimposed Non-Linear Measurements”. *IEEE Trans. Inf. Theory* 66.1 (2020), 453–477.
- [GK20] M. Genzel and C. Kipp. “Generic Error Bounds for the Generalized Lasso with Sub-Exponential Data”. Preprint arXiv:2004.05361. 2020.
- [GK18] M. Genzel and G. Kutyniok. “The Mismatch Principle: The Generalized Lasso Under Large Model Uncertainties”. Technical report, not intended for publication, arXiv:1808.06329. 2018.
- [GKM20] M. Genzel, G. Kutyniok, and M. März. “ ℓ^1 -Analysis Minimization and Generalized (Co-)Sparsity: When Does Recovery Succeed?” *Appl. Comput. Harmon. Anal.* (2020), in press, <https://doi.org/10.1016/j.acha.2020.01.002>.
- [GMS20] M. Genzel, M. März, and R. Seidel. “Compressed Sensing with 1D Total Variation: Breaking Sample Complexity Barriers via Non-Uniform Recovery”. Preprint arXiv:2001.09952. 2020.
- [GS20] M. Genzel and A. Stollenwerk. “Robust 1-bit compressed sensing via hinge loss minimization”. *Inf. Inference* 9.2 (2020), 361–422.
- [GMW18] L. Goldstein, S. Minsker, and X. Wei. “Structured Signal Recovery From Non-Linear and Heavy-Tailed Measurements”. *IEEE Trans. Inf. Theory* 64.8 (2018), 5513–5530.
- [GN98] R. M. Gray and D. L. Neuhoff. “Quantization”. *IEEE Trans. Inf. Theory* 44.6 (1998), 2325–2383.
- [GS93] R. M. Gray and T. G. Stockham. “Dithered quantizers”. *IEEE Trans. Inf. Theory* 39.3 (1993), 805–812.
- [JC17] L. Jacques and V. Cambareri. “Time for dithering: fast and quantized random embeddings via the restricted isometry property”. *Inf. Inference* 6.4 (2017), 441–476.
- [JLBB13] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk. “Robust 1-Bit Compressive Sensing via Binary Stable Embeddings of Sparse Vectors”. *IEEE Trans. Inf. Theory* 59.4 (2013), 2082–2102.
- [JLPY20] H. Jeong, X. Li, Y. Plan, and Ö. Yilmaz. “Sub-Gaussian Matrices on Sets: Optimal Tail Dependence and Applications”. Preprint arXiv:2001.10631. 2020.
- [JMPS19] H. C. Jung, J. Maly, L. Palzer, and A. Stollenwerk. “Quantized Compressed Sensing by Rectified Linear Units”. Preprint arXiv:1911.07816. 2019.
- [KSW16] K. Knudson, R. Saab, and R. Ward. “One-Bit Compressive Sensing With Norm Estimation”. *IEEE Trans. Inf. Theory* 62.5 (2016), 2748–2758.
- [LS20] Z. Liu and J. Scarlett. “The Generalized Lasso with Nonlinear Observations and Generative Priors”. Preprint arXiv:2006.12415. 2020.
- [MBKW20] M. März, C. Boyer, J. Kahn, and P. Weiss. “Sampling Rates for ℓ^1 -Synthesis”. Preprint arXiv:2004.07175. 2020.
- [Men16] S. Mendelson. “Upper bounds on product and multiplier empirical processes”. *Stoch. Proc. Appl.* 126.12 (2016), 3652–3680.
- [MPT07] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. “Reconstruction and Subgaussian Operators in Asymptotic Geometric Analysis”. *Geom. Funct. Anal.* 17.4 (2007), 1248–1282.
- [MJCDD16] A. Moshtaghpour, L. Jacques, V. Cambareri, K. Degraux, and C. De Vleeschouwer. “Consistent basis pursuit for signal and matrix estimates in quantized compressed sensing”. *IEEE Signal Process. Lett.* 23.1 (2016), 25–29.
- [ORS18] S. Oymak, B. Recht, and M. Soltanolkotabi. “Sharp Time-Data Tradeoffs for Linear Inverse Problems”. *IEEE Trans. Inf. Theory* 64.6 (2018), 4129–4158.
- [OR15] S. Oymak and B. Recht. “Near-Optimal Bounds for Binary Embeddings of Arbitrary Sets”. Preprint arXiv:1512.04433. 2015.
- [OS17] S. Oymak and M. Soltanolkotabi. “Fast and Reliable Parameter Estimation from Nonlinear Observations”. *SIAM J. Optim.* 27.4 (2017), 2276–2300.
- [PV13a] Y. Plan and R. Vershynin. “One-Bit Compressed Sensing by Linear Programming”. *Comm. Pure Appl. Math.* 66.8 (2013), 1275–1297.

- [PV13b] Y. Plan and R. Vershynin. “Robust 1-bit Compressed Sensing and Sparse Logistic Regression: A Convex Programming Approach”. *IEEE Trans. Inf. Theory* 59.1 (2013), 482–494.
- [PV14] Y. Plan and R. Vershynin. “Dimension Reduction by Random Hyperplane Tessellations”. *Discrete Comput. Geom.* 51.2 (2014), 438–461.
- [PV16] Y. Plan and R. Vershynin. “The Generalized Lasso With Non-Linear Observations”. *IEEE Trans. Inf. Theory* 62.3 (2016), 1528–1537.
- [PYY16] Y. Plan, R. Vershynin, and E. Yudovina. “High-dimensional estimation with geometric constraints”. *Inf. Inference* 6.1 (2016), 1–40.
- [ROV14] E. Richard, G. R. Obozinski, and J.-P. Vert. “Tight convex relaxations for sparse matrix factorization”. *Advances in Neural Information Processing Systems* 27. Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger. 2014, 3284–3292.
- [RV08] M. Rudelson and R. Vershynin. “On sparse reconstruction from Fourier and Gaussian measurements”. *Comm. Pure Appl. Math.* 61.8 (2008), 1025–1045.
- [SO20] Y. Sattar and S. Oymak. “Quickly Finding the Best Linear Model in High Dimensions via Projected Gradient Descent”. *IEEE Trans. Signal Process.* 68 (2020), 818–829.
- [Sto09] M. Stojnic. “Various thresholds for ℓ_1 -optimization in compressed sensing”. Preprint arXiv:0907.3666. 2009.
- [Tal14] M. Talagrand. *Upper and Lower Bounds for Stochastic Processes*. Vol. 3. Ergebnisse der Mathematik und ihrer Grenzgebiete. Springer Berlin Heidelberg, 2014.
- [TR17] C. Thrampoulidis and A. S. Rawat. “Lifting high-dimensional non-linear models with Gaussian regressors”. Preprint arXiv:1712.03638. 2017.
- [TAH15] C. Thrampoulidis, E. Abbasi, and B. Hassibi. “LASSO with Non-linear Measurements is Equivalent to One With Linear Measurements”. *Advances in Neural Information Processing Systems* 28. Ed. by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett. Curran Associates, 2015, 3402–3410.
- [TR20] C. Thrampoulidis and A. S. Rawat. “The Generalized Lasso for Sub-gaussian Measurements with Dithered Quantization”. *IEEE Trans. Inf. Theory* 66.4 (2020), 2487–2500.
- [Ver18] R. Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Vol. 47. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.
- [XJ18] C. Xu and L. Jacques. “Quantized Compressive Sensing with RIP Matrices: The Benefit of Dithering”. Preprint arXiv:1801.05870. 2018.
- [YBWL17] Z. Yang, K. Balasubramanian, Z. Wang, and H. Liu. “Estimating High-dimensional Non-Gaussian Multiple Index Models via Stein’s Lemma”. *Advances in Neural Information Processing Systems* 30. Ed. by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Curran Associates, 2017, 6097–6106.