Robust 1-Bit Compressed Sensing via Hinge Loss Minimization

Martin Genzel
Department of Mathematics
Technische Universität Berlin
E-mail: genzel@math.tu-berlin.de

Alexander Stollenwerk
Chair for Mathematics C (Analysis)
RWTH Aachen University
E-mail: stollenwerk@mathc.rwth-aachen.de

***Abstract*—We study the problem of estimating a structured high-dimensional signal $\boldsymbol{x}_0 \in \mathbb{R}^n$ from noisy 1-bit Gaussian measurements. Our recovery approach is based on a simple convex program which uses the hinge loss function as data fidelity term. While such a risk minimization strategy is typically applied in classification tasks, its capacity to estimate a specific signal vector is largely unexplored. In contrast to other popular loss functions considered in signal estimation, which are at least locally strongly convex, the hinge loss is just piecewise linear, so that its "curvature energy" is concentrated in a single point. It is therefore somewhat unexpected that we can still prove very similar types of recovery guarantees for the hinge loss estimator, even in the presence of strong noise. More specifically, our error bounds show that stable and robust reconstruction of $\boldsymbol{x}_0$ can be achieved with the optimal approximation rate $O(m^{-1/2})$ in terms of the number of measurements m. Moreover, we permit a wide class of structural assumptions on the ground truth signal, in the sense that $\boldsymbol{x}_0$ can belong to an arbitrary bounded convex set $K \subset \mathbb{R}^n$. For the proofs of our main results we invoke an adapted version of Mendelson's small ball method that allows us to establish a quadratic lower bound on the error of the first order Taylor approximation of the empirical hinge loss function.**

I. Motivation

We consider the problem of estimating an unknown *signal vector* $\boldsymbol{x}_0 \in \mathbb{R}^n$ from 1-bit observations of the form

$$y_i = f_i(\langle \boldsymbol{a}_i, \boldsymbol{x}_0\rangle) \in \{-1,+1\}, \quad i = 1,\dots,m, \tag{1}$$

where $\boldsymbol{a}_1,\dots,\boldsymbol{a}_m \in \mathbb{R}^n$ is a collection of known *measurement vectors* and $f_i \colon \mathbb{R} \to \{-1,+1\}$, $i = 1,\dots,m$, are binary-valued *quantization functions*. The number of samples m is typically much smaller than the ambient dimension n, so that the equation system of (1) is highly underdetermined. Such types of recovery tasks have recently caught increasing attention in various research areas, most importantly in the field of 1-*bit compressed sensing* [1], [2], [3]. The quantization step is motivated by real-world sensing schemes in which only a finite number of bits can be (digitally) processed during transmission. Let us emphasize that the quantizers f_i can be completely deterministic, e.g., $f_i = \operatorname{sign}$, but they could be also contaminated by noise in the form of random bit flips.

A large class of signal estimation methods can be formulated as an empirical risk minimization problem of the form

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \tfrac{1}{m}\sum_{i=1}^m \mathcal{L}(\langle \boldsymbol{a}_i, \boldsymbol{x}\rangle, y_i) \quad \text{subject to} \quad \boldsymbol{x} \in K, \tag{$P_{\mathcal{L},K}$}$$

where $\mathcal{L}\colon \mathbb{R}\times\mathbb{R}\to\mathbb{R}$ is a convex loss function that assesses how well the candidate model $\boldsymbol{a}_i \mapsto \langle \boldsymbol{a}_i, \boldsymbol{x}\rangle$ matches with the true outputs y_i, and $K \subset \mathbb{R}^n$ is a constraint set, which encodes the structural assumptions on $\boldsymbol{x}_0$ (e.g. sparsity). In this work, we will focus on a special instance of ($P_{\mathcal{L},K}$) that is based on the so-called *hinge loss* given by $\mathcal{L}^{\text{hng}}(v) := [1-v]_+ := \max\{0, 1-v\}$ for $v \in \mathbb{R}$. Using $\mathcal{L}(v,v') := \mathcal{L}^{\text{hng}}(v\cdot v')$ as loss function, the program of ($P_{\mathcal{L},K}$) reads as follows:

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \tfrac{1}{m}\sum_{i=1}^m [1 - y_i\langle \boldsymbol{a}_i, \boldsymbol{x}\rangle]_+ \quad \text{subject to} \quad \boldsymbol{x} \in K. \tag{$P_{\mathcal{L}^{\text{hng}},K}$}$$

This estimator is specifically tailored to deal with binary observations: Intuitively, by minimizing the objective functional of ($P_{\mathcal{L}^{\text{hng}},K}$), one tries to select $\boldsymbol{x} \in K$ in such a way that $\operatorname{sign}(\langle \boldsymbol{a}_i, \boldsymbol{x}\rangle)$ equals $y_i \in \{-1,+1\}$ for as many samples as possible. In particular, a solution $\hat{\boldsymbol{x}}$ to ($P_{\mathcal{L}^{\text{hng}},K}$) yields a good predictor $\boldsymbol{a}_i \mapsto \operatorname{sign}(\langle \boldsymbol{a}_i, \hat{\boldsymbol{x}}\rangle)$ of the true outputs y_i, based on the assumption that this can be achieved with a vector from K. While this simple heuristic explains the success of hinge loss minimization in classification problems, the performance of ($P_{\mathcal{L}^{\text{hng}},K}$) at signal estimation tasks is only poorly understood. Compared to reliable prediction, successful signal estimation usually relies on relatively strong model assumptions which ensure that one actually retrieves the ground truth signal $\boldsymbol{x}_0$ and not just *any* good predictor. Our goal is therefore to establish theoretical recovery guarantees for ($P_{\mathcal{L}^{\text{hng}},K}$) under the hypothesis of (1) with *Gaussian* measurement vectors.

A. Main Contributions

Recent results, e.g. [4], [5], show that for a large class of loss functions $\mathcal{L}$, the estimator ($P_{\mathcal{L},K}$) is capable of reconstructing structured signals $\boldsymbol{x}_0 \in K$ from quite general *non-linear* (including 1-bit) measurements. While this includes popular choices of $\mathcal{L}$, such as the logistic loss or the square loss, the hinge loss does not meet known sufficient conditions which ensure successful signal estimation. In this work, we show that 1-bit compressed sensing via hinge loss minimization is feasible for large classes of signal models K and 1-bit observation schemes (1). In particular, our guarantees significantly improve a recent result from Kolleck and Vybíral [6], whose analysis of the hinge loss estimator is limited to ℓ^1-constraints and a far more restrictive noise pattern. Moreover, the error bounds in [6] do only achieve an approximation rate of $O(m^{-1/4})$, while Theorem 1 and Theorem 3 below exhibit the (optimal) rate of $O(m^{-1/2})$. Our improvements are due to a substantially different *localized* analysis, in which the "quadratic" part of the hinge loss is explicitly taken into account.

II. Main Results

This part presents our main theoretical findings on signal estimation via hinge loss minimization. Let us begin by giving a precise definition of the observation model that was informally introduced in (1):

Assumption 1 (Measurement Model)**.** *Let $f : \mathbb{R} \to \{-1, +1\}$ be a (random) quantization function and let $\boldsymbol{a} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I}_n)$ be a standard Gaussian random vector which is independent of f. We consider a* noisy 1-bit Gaussian measurement model *of the form*

$$y := f(\langle \boldsymbol{a}, \boldsymbol{x}_0 \rangle) \in \{-1, +1\}$$

where $\boldsymbol{x}_0 \in \mathbb{R}^n$ is the (unknown) ground truth signal. *Each of the m samples $\{(\boldsymbol{a}_i, y_i)\}_{i\in[m]} \subset \mathbb{R}^n \times \{-1, +1\}$ is then drawn as an independent copy from the random pair $(\boldsymbol{a}, y)$. Consequently, the binary observations are given by*

$$y_i = f_i(\langle \boldsymbol{a}_i, \boldsymbol{x}_0 \rangle), \quad i = 1, \dots, m,$$

where f_i is an independent copy of f.

The prototypical example of a 1-bit quantizer is the sign-function, that is, $f = \operatorname{sign}$. We refer to this (noiseless) observation scheme as the *perfect 1-bit model*. Since all information on the magnitude of $\boldsymbol{x}_0$ is lost in this case, we will additionally assume that the signal vector $\boldsymbol{x}_0$ is normalized. Moreover, in many scenarios of interest, we have some prior knowledge on the signal's structure. The hinge loss estimator $(P_{\mathcal{L}^{\text{hng}},K})$ encodes such structural assumptions by means of a constraint set $K \subset \mathbb{R}^n$. Hence, we supplement Assumption 1 with the following signal model:

Assumption 2 (Signal Model)**.** *We assume that $\|\boldsymbol{x}_0\|_2 = 1$ and $\boldsymbol{x}_0 \in K$ for a certain subset $K \subset \mathbb{R}^n$, which is called the* signal set. *Furthermore, we require that K is convex, bounded, and $\mathbf{0} \in K$.*

The most prominent signal structure in compressed sensing is *sparsity*. A signal vector $\boldsymbol{x}_0 \in \mathbb{R}^n$ is called sparse, if $\|\boldsymbol{x}_0\|_0 \leq s$ for some $s \ll n$. The set of all unit-norm s-sparse vectors is however not convex, so that Assumption 2 is not fulfilled. Nevertheless, the Cauchy-Schwarz inequality implies that both $K = \sqrt{s}B_1^n$ and $K = \sqrt{s}B_1^n \cap B_2^n$ may serve as admissible convex relaxations, which meet the conditions of Assumption 2.

Given prior information on the structure of the target vector, a key issue in signal estimation concerns the design of recovery procedures that effectively exploit this information and ensure recovery from a number of measurements m that scales as the complexity of this structure. If the signal vector's structure is encoded by means of a signal set K, its *Gaussian width* has turned out to be a useful complexity measure:

Definition II.1 The *Gaussian width* of a bounded set $K \subset \mathbb{R}^n$ is defined as

$$w(K) := \mathbb{E}[\sup_{\boldsymbol{x}\in K} \langle \boldsymbol{g}, \boldsymbol{x} \rangle],$$

where $\boldsymbol{g} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I}_n)$ denotes a standard Gaussian random vector.

Many results — including ours below — show that the square of the Gaussian width often determines the (minimal) number of samples to ensure recovery via convex optimization [7], [8]. Since the Gaussian widths of $\sqrt{s}B_1^n \cap B_2^n$ and $\sqrt{s}B_1^n$ are bounded as follows (see [9, Sec. 2 and 3])

$$w(\sqrt{s}B_1^n) \lesssim \sqrt{s\log(n)}, \; w(\sqrt{s}B_1^n \cap B_2^n) \lesssim \sqrt{s\log(\tfrac{2n}{s})},$$

this implies that (approximately) s-sparse vectors are efficiently estimated if the number of measurements m is larger than s multiplied by a logarithmic factor in the ambient dimension. The second (localized) complexity measure that we need in order to formulate our recovery results is the so-called *conic effective dimension* [8]:

Definition II.2 (Conic Effective Dimension) The *conic effective dimension* (or *statistical dimension*) of a subset $K \subset \mathbb{R}^n$ in $\boldsymbol{x}_0$ is defined as

$$d_0(K - \boldsymbol{x}_0) := w(\mathcal{C}(K, \boldsymbol{x}_0) \cap B_2^n)^2,$$

where

$$\mathcal{C}(K, \boldsymbol{x}_0) := \{\tau(\boldsymbol{x} - \boldsymbol{x}_0) \mid \boldsymbol{x} \in K, \; \tau \geq 0\}$$

is the *descent cone* of K at $\boldsymbol{x}_0$.

The conic effective dimension is a complexity measure that, geometrically speaking, measures the size (narrowness) of the cone generated by $K - \boldsymbol{x}_0$. As an example, if $\boldsymbol{x}_0$ is s-sparse and lies on the boundary of an ℓ^1-ball, then (see [7, Prop. 3.10])

$$d_0(\|\boldsymbol{x}_0\|_1 B_1^n - \boldsymbol{x}_0) \lesssim s\log(\tfrac{2n}{s}).$$

Before we state our main results, let us briefly outline the general idea behind signal estimation via empirical risk minimization. By the law of large numbers, we expect that for m large enough, a solution $\hat{\boldsymbol{x}}$ to $(P_{\mathcal{L}^{\text{hng}},K})$, that is, a minimizer of the *empirical risk function*

$$\bar{\mathcal{R}}(\boldsymbol{x}) := \tfrac{1}{m}\sum_{i=1}^{m} \mathcal{L}^{\text{hng}}(y_i\langle \boldsymbol{a}_i, \boldsymbol{x} \rangle) = \tfrac{1}{m}\sum_{i=1}^{m} [1 - y_i\langle \boldsymbol{a}_i, \boldsymbol{x} \rangle]_+$$

on the set K, will be close to a minimizer $\boldsymbol{x}^*$ of the *expected risk function* $\mathcal{R}(\boldsymbol{x}) := \mathbb{E}[\bar{\mathcal{R}}(\boldsymbol{x})] = \mathbb{E}[\mathcal{L}^{\text{hng}}(y\langle \boldsymbol{a}, \boldsymbol{x} \rangle)]$ on K. Hence, in order for $\hat{\boldsymbol{x}}/\|\hat{\boldsymbol{x}}\|_2$ to be close to $\boldsymbol{x}_0 \in \mathbb{S}^{n-1}$, the vector $\boldsymbol{x}^*/\|\boldsymbol{x}^*\|_2$ has to be close to $\boldsymbol{x}_0$ as well. The last issue strongly depends on whether or not K is contained in B_2^n: Under a mild condition on the 1-bit quantizer f in Assumption 1, if $K \subset B_2^n$, then there exists a minimizer of the expected risk function on K which is contained in the span of $\boldsymbol{x}_0$, see Lemma 1. In contrast, even for the choice $f = \operatorname{sign}$, if K is not contained in B_2^n, then one generally cannot deduce that $\boldsymbol{x}^*/\|\boldsymbol{x}^*\|_2$ is close to $\boldsymbol{x}_0$. As a consequence, hinge loss minimization using the estimator $(P_{\mathcal{L}^{\text{hng}},K})$ might fail if $K \not\subset B_2^n$. The key observation is that by instead considering a minimizer $\boldsymbol{x}^*$ of the expected risk function on the larger set μK, where $\mu \geq 1$, then $\boldsymbol{x}^*/\|\boldsymbol{x}^*\|_2$ will be close to $\boldsymbol{x}_0$ for μ large enough, see Proposition 1. Hence, for general convex sets K we will instead consider a tunable version of the estimator $(P_{\mathcal{L}^{\text{hng}},K})$, see Definition II.3. In view of these structural differences, we sort our recovery results according to whether or not K is contained in B_2^n.

A. Recovery in Subsets of the Unit Ball

The first lemma shows that if $K \subset B_2^n$, then the expected risk function attains its minimum on the span of $\boldsymbol{x}_0$.

Lemma 1. *Let $g \sim \mathcal{N}(0,1)$ be a standard Gaussian random variable. Moreover, let $\mu \in [0,1]$ be a minimizer of*

$$\min_{s\in[0,1]} \mathbb{E}[\mathcal{L}^{hng}(sf(g)g)], \tag{2}$$

where $f\colon \mathbb{R} \to \{-1,+1\}$ is the 1-bit quantizer from Assumption 1. Assuming that $\mathbb{E}[f(g)g] > 0$, we have that $\mu > 0$ and $\mu\boldsymbol{x}_0 \in K$ satisfies

$$\mathcal{R}(\mu\boldsymbol{x}_0) = \min_{\boldsymbol{x}\in K} \mathcal{R}(\boldsymbol{x}).$$

Let us emphasize that $\mathbb{E}[f(g)g] > 0$ is a reasonable assumption because it ensures that the linear measurement $g = \langle \boldsymbol{a}, \boldsymbol{x}_0\rangle$ and the output variable $y = f(g)$ are positively correlated. We need a second mild condition on the quantizer f in order to formulate our main recovery result Theorem 1:

Assumption 3 (Correlation Conditions)**.** *Let $g \sim \mathcal{N}(0,1)$ be a standard Gaussian random variable and let $f\colon \mathbb{R} \to \{-1,+1\}$ be the 1-bit quantizer from Assumption 1. We assume that the following two model conditions hold true:*

1) $\lambda := \lambda_f := \mathbb{E}[f(g)g] > 0$,
2) $\mathbb{E}[f(g)\,\text{sign}(g) \mid |g|] \geq 0$ *(a.s.)*.

We call λ the correlation parameter *of the quantizer f.*

Next, we state our main recovery result for signal sets K that are contained in the Euclidean unit ball.

Theorem 1 (Signal Recovery in Unit Ball)**.** *Let the model conditions of Assumption 1, 2, and 3 be satisfied, assume that $K \subset B_2^n$, and let μ be defined according to* (2)*. For every $t \in (0,\mu)$ and $\eta \in (0,\frac{1}{2})$, the following holds true with probability at least $1-\eta$: If the number of samples obeys*

$$m \gtrsim \lambda^{-2}\cdot t^{-2}\cdot \max\{d_0(K-\mu\boldsymbol{x}_0), \log(\eta^{-1})\},$$

then any minimizer $\hat{\boldsymbol{x}}$ of $(P_{\mathcal{L}^{\text{hng}},K})$ satisfies

$$\left\|\boldsymbol{x}_0 - \frac{\hat{\boldsymbol{x}}}{\|\hat{\boldsymbol{x}}\|_2}\right\|_2 \leq \frac{t}{\mu} \lesssim t\cdot\sqrt{\log(\lambda^{-1})}\ .$$

A remarkable feature of Theorem 1 is that the impact of the underlying 1-bit measurement model is completely controlled by the correlation parameter λ. Since λ can be regarded as a constant scaling factor, recovery via $(P_{\mathcal{L}^{\text{hng}},K})$ is still possible when the specific output rule is unknown and the signal-to-noise ratio is very low. As an example, let us consider the case where independent sign flips corrupt the quantization process. Here, the 1-bit observations are given by

$$y_i = \varepsilon_i \cdot \text{sign}(\langle \boldsymbol{a}_i, \boldsymbol{x}_0\rangle), \quad i = 1,\dots,m,$$

where ε_i are independent copies of a Bernoulli random variable $\varepsilon \in \{-1,+1\}$ with $\mathbb{P}[\varepsilon = 1] = p > \frac{1}{2}$. This measurement model is an instance of Assumption 1 for the 1-bit quantizer

$$f(v) = f_p(v) := \varepsilon\cdot\text{sign}(v).$$

Notice that $p = 1$ corresponds to perfect (noiseless) 1-bit observations. It is straightforward to verify that this noisy 1-bit model satisfies Assumption 3 with correlation parameter $\lambda_{f_p} = (2p-1)\sqrt{\frac{2}{\pi}}$. Consequently, Theorem 1 shows that signal recovery succeeds if the number of measurements m satisfies

$$m \gtrsim \tfrac{1}{(2p-1)^2}\cdot C_{t,K},$$

where the constant $C_{t,K} > 0$ hides the dependence on the oversampling factor t and the signal complexity. Hence, signal recovery is still feasible in the presence of strong noise where $p \approx \frac{1}{2}$.

B. Recovery in General Convex Sets

Sometimes, it can be computationally appealing to drop the unit-ball assumption and to allow for "larger" convex signal sets. A common example is an ℓ^1-penalty, for which $(P_{\mathcal{L}^{\text{hng}},K})$ can be reformulated as a linear program (cf. [6, Sec. VI.A]). This motivates us to investigate the recovery performance of hinge loss minimization under arbitrary convex constraints. We will make the following model assumptions throughout this subsection:

Assumption 4 (General Signal Sets)**.** *Let $\boldsymbol{x}_0 \in \mathbb{S}^{n-1}$ be a unit vector in $\mathbb{R}^n$ and let $\boldsymbol{a} \sim \mathcal{N}(\boldsymbol{0},\boldsymbol{I}_n)$ be a standard Gaussian. We consider* perfect 1-bit Gaussian measurements

$$y_i = \text{sign}(\langle \boldsymbol{a}_i, \boldsymbol{x}_0\rangle), \quad i = 1,\dots,m.$$

Further, we assume that $\boldsymbol{x}_0 \in K$ for a signal set $K \subset \mathbb{R}^n$ which is convex, bounded, and closed.

Under Assumption 4, a (normalized) minimizer of the expected risk function might no longer be close to $\boldsymbol{x}_0$. The picture changes completely if we upscale the signal set K:

Proposition 1. *Let Assumption 4 be satisfied and assume that $\mu \gtrsim 1$. Then every expected risk minimizer $\boldsymbol{x}^*$ on μK (i.e., $\mathcal{R}(\boldsymbol{x}^*) = \min_{\boldsymbol{x}\in\mu K}\mathcal{R}(\boldsymbol{x})$) satisfies*

$$\left\|\boldsymbol{x}_0 - \frac{\boldsymbol{x}^*}{\|\boldsymbol{x}^*\|_2}\right\|_2 \lesssim \frac{1}{\mu}\ .$$

This motivates us to introduce an adapted version of $(P_{\mathcal{L}^{\text{hng}},K})$ that allows us to rescale the signal set:

Definition II.3 (Scalable Hinge Loss Minimization) Let Assumption 4 hold true and let $\mu > 0$ be a fixed *scaling* *parameter*. The estimator $\hat{\boldsymbol{x}} \in \mathbb{R}^n$ is defined as a solution of the convex program

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \tfrac{1}{m}\sum_{i=1}^{m} [1 - y_i\langle \boldsymbol{a}_i, \boldsymbol{x}\rangle]_+ \quad \text{subject to} \quad \boldsymbol{x} \in \mu K. \tag{$P_{\mathcal{L}^{\text{hng}},\mu K}$}$$

The next result shows that scalable hinge loss minimization is an efficient reconstruction procedure under Assumption 4:

Theorem 2 (Signal Recovery in Convex Sets)**.** *Let the model conditions of Assumption 4 be satisfied. For every fixed $\mu > 0$ and $\eta \in (0, \frac{1}{2})$, the following holds true with probability at least $1 - \eta$: If $\mu \gtrsim 1$ and the number of samples obeys*

$$m \gtrsim \mu^4 \cdot \max\{d_0(K - \boldsymbol{x}_0), \log(\eta^{-1})\}, \tag{3}$$

then any minimizer $\hat{\boldsymbol{x}}$ of $(P_{\mathcal{L}^{\text{hng}},\mu K})$ satisfies

$$\left\| \boldsymbol{x}_0 - \frac{\hat{\boldsymbol{x}}}{\|\hat{\boldsymbol{x}}\|_2} \right\|_2 \lesssim \frac{1}{\mu}\ .$$

The same assertion holds true if (3) *is replaced by*

$$m \gtrsim \mu^4 \cdot \max\{w(K)^2, \log(\eta^{-1})\}.$$

A downside of Theorem 2 is the relatively slow error decay of $O(m^{-1/4})$. Our third main result, Theorem 3, shows that the factor of μ^4 in condition (3) can be replaced by $\mu^2 \cdot t_0^2$, where t_0 is an additional geometric parameter that depends on $\boldsymbol{x}_0$ and K. Formally, it is defined by

$$t_0 := \max\{1, \operatorname{rad}((\partial(\mu K) \cap \operatorname{Cyl}(\boldsymbol{x}_0, \mu)) - \mu\boldsymbol{x}_0)\} \tag{4}$$

where $\operatorname{Cyl}(\boldsymbol{x}_0, \mu)$ denotes the following cylindrical tube around $\operatorname{span}\{\boldsymbol{x}_0\}$:

$$\operatorname{Cyl}(\boldsymbol{x}_0, \mu) := \{\boldsymbol{x} \in \mathbb{R}^n \mid \|\boldsymbol{x} - \langle \boldsymbol{x}, \boldsymbol{x}_0\rangle \boldsymbol{x}_0\|_2 \leq 1, \langle \boldsymbol{x}, \boldsymbol{x}_0\rangle \geq \tfrac{\mu}{2}\}.$$

Intuitively, if μ is sufficiently large, the boundary $\partial(\mu K)$ does only intersect with the side of $\operatorname{Cyl}(\boldsymbol{x}_0, \mu)$, which has diameter 2. The value of t_0 then becomes (almost) independent of μ and is determined by the "local" geometry of K in a neighborhood of $\boldsymbol{x}_0$ (of size $1/\mu$). In this situation, the sampling rate only scales quadratically in μ:

Theorem 3 (Signal Recovery in Convex Sets – Local Version)**.** *Let the model conditions of Assumption 4 be satisfied and let t_0 be defined according to* (4)*. For every fixed $\mu > 0$ and $\eta \in (0, \frac{1}{2})$, the following holds true with probability at least $1 - \eta$: If $\mu \gtrsim t_0$ and the number of samples obeys*

$$m \gtrsim \mu^2 \cdot t_0^2 \cdot \max\{d_0(K - \boldsymbol{x}_0), \log(\eta^{-1})\},$$

then any minimizer $\hat{\boldsymbol{x}}$ of $(P_{\mathcal{L}^{\text{hng}},\mu K})$ satisfies

$$\left\| \boldsymbol{x}_0 - \frac{\hat{\boldsymbol{x}}}{\|\hat{\boldsymbol{x}}\|_2} \right\|_2 \lesssim \frac{1}{\mu}\ .$$

Without any further assumptions on the geometric arrangement of K and $\boldsymbol{x}_0$, it is difficult to make precise statements about the order of t_0. For example, if the boundary of K is almost orthogonal to $\operatorname{span}\{\boldsymbol{x}_0\}$ in a small neighborhood of $\boldsymbol{x}_0$, we can expect that $t_0 \approx 1$. But as ∂K gets more "tangent" to $\operatorname{span}\{\boldsymbol{x}_0\}$, t_0 may become significantly larger. However, as long as t_0 is considered as a (possibly large) signal-dependent parameter, we can always achieve the optimal rate of $O(m^{-1/2})$.

III. Proof Strategy

In the following, we give a proof sketch of our main recovery results Theorem 1–3. For a complete proof, see [10, Section 6]. Our main strategy is to show that, with high probability, any minimizer $\hat{\boldsymbol{x}}$ of the (scalable) hinge loss minimization program resides in a certain neighborhood of $\mu\boldsymbol{x}_0$ for a scaling parameter $\mu > 0$. If this neighborhood is chosen appropriately (depending on whether or not $K \subset B_n^2$), then the projection of $\hat{\boldsymbol{x}}$ onto the Euclidean unit sphere will be close to $\boldsymbol{x}_0$. In order to show that $\hat{\boldsymbol{x}}$ lies in a neighborhood of $\mu\boldsymbol{x}_0$, it suffices to show that the convex *excess risk*

$$\mathcal{E}(\boldsymbol{x}) := \bar{\mathcal{R}}(\boldsymbol{x}) - \bar{\mathcal{R}}(\mu\boldsymbol{x}_0), \quad \boldsymbol{x} \in \mathbb{R}^n,$$

is positive on the boundary of this neighborhood. Such a localization argument is widely used in estimation theory and statistical learning. In order to show positivity of the excess risk, we consider the first order Taylor expansion of $\boldsymbol{x} \mapsto \bar{\mathcal{R}}(\boldsymbol{x})$ at $\mu\boldsymbol{x}_0$. The approximation error is then given by

$$\mathcal{Q}(\boldsymbol{x}, \mu\boldsymbol{x}_0) := \bar{\mathcal{R}}(\boldsymbol{x}) - \underbrace{\bar{\mathcal{R}}(\mu\boldsymbol{x}_0) - \tfrac{1}{m}\sum_{i=1}^{m} z_i\langle \boldsymbol{a}_i, \boldsymbol{x} - \mu\boldsymbol{x}_0\rangle}_{=:\mathcal{M}(\boldsymbol{x},\mu\boldsymbol{x}_0)},$$

where $\mathcal{M}(\cdot, \mu\boldsymbol{x}_0)$ is the "linearization" of $\bar{\mathcal{R}}(\cdot)$ at $\mu\boldsymbol{x}_0$ with

$$z_i := y_i \cdot [\mathcal{L}^{\text{hng}}]'(y_i\langle \boldsymbol{a}_i, \mu\boldsymbol{x}_0\rangle) = -y_i \cdot \chi_{(-\infty,1]}(y_i\langle \boldsymbol{a}_i, \mu\boldsymbol{x}_0\rangle).$$

By convexity of $\bar{\mathcal{R}}(\cdot)$ the "quadratic" term $\mathcal{Q}(\cdot, \mu\boldsymbol{x}_0)$ is always non-negative. Hence, in order to achieve

$$\mathcal{E}(\boldsymbol{x}) = \bar{\mathcal{R}}(\boldsymbol{x}) - \bar{\mathcal{R}}(\mu\boldsymbol{x}_0) = \mathcal{M}(\boldsymbol{x}, \mu\boldsymbol{x}_0) + \mathcal{Q}(\boldsymbol{x}, \mu\boldsymbol{x}_0) > 0$$

for all $\boldsymbol{x}$ in a fixed subset of $\mathbb{R}^n$, it suffices to show that $\mathcal{Q}(\cdot, \mu\boldsymbol{x}_0)$ uniformly dominates $\mathcal{M}(\cdot, \mu\boldsymbol{x}_0)$ on that specific set. The term $\mathcal{M}(\cdot, \mu\boldsymbol{x}_0)$ can be handled by a recent result of Mendelson [11, Thm. 4.4] which concerns the uniform deviation of multiplier empirical processes from their mean. After bounding the quadratic term $\mathcal{Q}(\boldsymbol{x}, \mu\boldsymbol{x}_0)$ from below by a simplified non-negative empirical process, this process is again controlled using an adaption of Tropp's version of *Mendelson's small ball method* in [12, Prop. 5.1].

Acknowledgements

The authors would like to thank Sjoerd Dirksen for initiating this project and for many fruitful discussions. M.G. is supported by the Bundesministerium für Bildung und Forschung (BMBF) through the Berliner Zentrum for Machine Learning (BZML), Project TP4. A.S acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG) within the priority program SPP 1798 Compressed Sensing in Information Processing through the project Quantized Compressive Spectrum Sensing (QuaCoSS).

REFERENCES

[1] P. T. Boufounos and R. G. Baraniuk, "1-bit compressive sensing," in *Proceedings of the 42nd Annual Conference on Information Sciences and Systems (CISS 2008)*, 2008, pp. 16–21.

[2] 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*, vol. 59, no. 4, pp. 2082–2102, 2013.

[3] Y. Plan and R. Vershynin, "One-bit compressed sensing by linear programming," *Comm. Pure Appl. Math.*, vol. 66, no. 8, pp. 1275–1297, 2013.

[4] ——, "The generalized Lasso with non-linear observations," *IEEE Trans. Inf. Theory*, vol. 62, no. 3, pp. 1528–1537, 2016.

[5] M. Genzel, "High-dimensional estimation of structured signals from non-linear observations with general convex loss functions," *IEEE Trans. Inf. Theory*, vol. 63, no. 3, pp. 1601–1619, 2017.

[6] A. Kolleck and J. Vybíral, "Non-asymptotic analysis of ℓ_1-norm support vector machines," *IEEE Trans. Inf. Theory*, vol. 63, no. 9, pp. 5461–5476, 2017.

[7] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The convex geometry of linear inverse problems," *Found. Comput. Math.*, vol. 12, no. 6, pp. 805–849, 2012.

[8] 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*, vol. 3, no. 3, pp. 224–294, 2014.

[9] Y. Plan and R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach," *IEEE Trans. Inf. Theory*, vol. 59, no. 1, pp. 482–494, 2013.

[10] M. Genzel and A. Stollenwerk, "Robust 1-bit compressed sensing via hinge loss minimization," *CoRR*, vol. abs/1804.04846, 2018.

[11] S. Mendelson, "Upper bounds on product and multiplier empirical processes," *Stoch. Process. Appl.*, vol. 126, no. 12, pp. 3652–3680, 2016.

[12] J. A. Tropp, "Convex recovery of a structured signal from independent random linear measurements," in *Sampling Theory, a Renaissance*, ser. Applied and Numerical Harmonic Analysis, G. E. Pfander, Ed. Birkhäuser, 2015, pp. 67–101.