



Sjoerd Dirksen · Shahar Mendelson

Non-Gaussian hyperplane tessellations and robust one-bit compressed sensing

Received August 13, 2018

Abstract. We show that a tessellation generated by a small number of random affine hyperplanes can be used to approximate Euclidean distances between any two points in an arbitrary bounded set T , where the random hyperplanes are generated by subgaussian or heavy-tailed normal vectors and uniformly distributed shifts. The number of hyperplanes needed for constructing such tessellations is determined by natural metric complexity measures of the set T and the wanted approximation error. In comparison, previous results in this direction were restricted to Gaussian hyperplane tessellations of subsets of the Euclidean unit sphere.

As an application, we obtain new reconstruction results in memoryless one-bit compressed sensing with non-Gaussian measurement matrices: by quantizing at uniformly distributed thresholds, it is possible to accurately reconstruct low-complexity signals from a small number of one-bit quantized measurements, even if the measurement vectors are drawn from a heavy-tailed distribution. These reconstruction results are uniform in nature and robust in the presence of pre-quantization noise on the analog measurements as well as adversarial bit corruptions in the quantization process. Moreover, if the measurement matrix is subgaussian then accurate recovery can be achieved via a convex program.

Keywords. Hyperplane tessellations, compressed sensing, quantization, empirical processes

1. Introduction

In what follows we study the following geometric question: can distances between points in a given set $T \subset \mathbb{R}^n$ be accurately encoded using a small number of random hyperplanes? To formulate the question more precisely, let $H_{X_i, \tau_i} = \{x \in \mathbb{R}^n : \langle X_i, x \rangle + \tau_i = 0\}$, $i = 1, \dots, m$, be a collection of affine hyperplanes with normal vectors X_i and shift parameters τ_i . These hyperplanes tessellate the set T into (at most) 2^m cells and, for any $x \in T$, the bit string $(\text{sign}(\langle X_i, x \rangle + \tau_i))_{i=1}^m \in \{-1, 1\}^m$ encodes the cell in which x is located (see Figures 1 and 2). Moreover, for any two points $x, y \in T$, the normalized

S. Dirksen: RWTH Aachen University,
Pontdriesch 10, 52062 Aachen, Germany; e-mail: dirksen@mathc.rwth-aachen.de

S. Mendelson: Mathematical Sciences Institute, The Australian National University,
ACT 2600, Canberra, Australia; e-mail: shahar.mendelson@anu.edu.au

Mathematics Subject Classification (2020): Primary 60D05; Secondary 60B20, 94A12




Fig. 1. An illustration of the hyperplane cut generated by the vector X_1 (and shift parameter 0). The homogeneous hyperplane H_{X_1} divides \mathbb{R}^n into two parts, a “+” and a “-” side. The red and green points are assigned the bit 1, the orange point is assigned -1 .




Fig. 2. The homogeneous hyperplanes H_{X_1} and H_{X_2} divide \mathbb{R}^n into four parts. The red, green, and orange points are assigned the bit sequences $\{1, 1\}$, $\{1, -1\}$ and $\{-1, -1\}$, respectively.

Hamming distance between their bit strings,

$$\frac{1}{m} |\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}|, \tag{1.1}$$

counts the fraction of hyperplanes separating x and y . In what follows the goal is to quantify the number of random hyperplanes that suffice to ensure that (1.1) approximates the distance between any two points in T that are not ‘too close’.

A beautiful result due to Plan and Vershynin [22] essentially solves this question for subsets of the Euclidean unit sphere with respect to the geodesic distance, using homogeneous Gaussian hyperplanes (i.e., $\tau_i = 0$ for all i). They showed that if $T \subset S^{n-1}$ and the normal vectors X_1, \dots, X_m are independent standard Gaussian vectors, then with

probability at least $1 - 2e^{-cm\rho^2}$, for all $x, y \in T$,

$$d_{S^{n-1}}(x, y) - \rho \leq \frac{1}{m} |\{i : \text{sign}(\langle X_i, x \rangle) \neq \text{sign}(\langle X_i, y \rangle)\}| \leq d_{S^{n-1}}(x, y) + \rho, \quad (1.2)$$

provided that $m \gtrsim \rho^{-6} \ell_*^2(T)$; here

$$\ell_*(T) := \mathbb{E} \sup_{x \in T} |\langle G, x \rangle|$$

and G is the standard Gaussian random vector in \mathbb{R}^n . Thus, $\ell_*(T)$ is the Gaussian mean-width of T —a natural geometric parameter that is of central importance in geometry (e.g. in Dvoretzky type theorems, see for instance [2]) and in statistics, where it is used to capture the difficulty of prediction problems.

It follows from (1.2) that if x and y are ‘far-enough apart’, then the fraction of homogeneous Gaussian hyperplanes that separate them concentrates sharply around their geodesic distance.

As far as random homogeneous Gaussian tessellations of $T \subset S^{n-1}$ are concerned, it was conjectured in [22] that $m \simeq \rho^{-2} \ell_*^2(T)$ is necessary and sufficient for (1.2) to hold. The best known sufficient condition for an arbitrary $T \subset S^{n-1}$ is $m \gtrsim \rho^{-4} \ell_*^2(T)$, established in [19], while for certain ‘simple’ subsets of the Euclidean sphere (e.g., if T is the intersection of a subspace and the sphere) $m \gtrsim \rho^{-2} \ell_*^2(T)$ is known to be sufficient [19, 22].

It is natural to ask whether approximating distances via random tessellations is possible in more general situations, most notably, using other distributions for generating the normal vectors rather than the standard Gaussian distribution, and considering sets T that need not be subsets of S^{n-1} . As it happens, these are not only natural extensions but, in fact, are of extreme importance in signal processing—specifically, when studying signal reconstruction problems from quantized measurements. The connection between the extended version of the random tessellation problem and signal recovery is explained in detail in Section 1.1.

Unfortunately, it is clear that the two extensions one is interested in are not possible when considering tessellations generated by homogeneous hyperplanes. First of all, it is impossible to separate points lying on a ray originating from 0 using a homogeneous hyperplane. And second, it is easy to find very natural distributions for which (1.2) is false. As an extreme case, observe that there are vectors in S^{n-1} that are far apart but still cannot be separated using H_{X_i} if X_1, \dots, X_m are selected according to the uniform distribution on $\{-1, 1\}^n$. In fact, the points cannot be separated even if one uses *all* possible hyperplanes generated by points in $\{-1, 1\}^n$.

A possible solution to both problems stems in a phenomenon that appears in engineering literature: there is extensive experimental evidence that signal recovery from quantized measurements improves substantially if one adds appropriate ‘noise’ to the measurements before quantizing. The operation of adding noise before quantization, which was first proposed in [23], is called *dithering* (see also the survey [12]).

In the context of random tessellations, the geometric interpretation of dithering is adding random parallel shifts to the hyperplanes. We show that adding such random shifts

allows one to address the two problems, and as a result, random tessellations of arbitrary sets T that are generated by rather general distributions can be used to approximate distances in T . Moreover, the reason why dithering is such an effective method in signal recovery problems becomes clear thanks to the analysis presented in what follows (see Section 1.1 for more details).

To formulate the main results of this article, consider i.i.d. shifts τ_i that are uniformly distributed in $[-\lambda, \lambda]$ for a well chosen λ , let RB_2^n be the Euclidean ball of radius R , and let $T \subset RB_2^n$. Set X to be a random vector in \mathbb{R}^n and let X_1, \dots, X_m be independent copies of X that are also independent of $(\tau_i)_{i=1}^m$.

Although the method introduced in what follows can be used in other situations (see in particular Remark 1.15), the focus here is on two scenarios.

The first scenario is called *the L -subgaussian scenario*, in which X is isotropic, symmetric, and L -subgaussian.¹ The following result is a special case of Theorem 2.3 below, and to formulate it, denote by $\text{conv}(T)$ the convex hull of the set T .

Theorem 1.1. *Set*

$$d(x, y) = \frac{1}{m} |\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}|.$$

There exist constants c_0, \dots, c_4 depending only on L such that the following holds. Fix $0 < \rho < R$. If $T \subset RB_2^n$, $\lambda = c_0R$ and

$$m \geq c_1 \frac{R \log(eR/\rho)}{\rho^3} \ell_*^2(T),$$

then with probability at least $1 - 2 \exp(-c_2 m \rho / R)$, for any $x, y \in \text{conv}(T)$ such that $\|x - y\|_2 \geq \rho$, one has

$$c_3 \frac{\|x - y\|_2}{R} \leq d(x, y) \leq c_4 \sqrt{\log(eR/\rho)} \cdot \frac{\|x - y\|_2}{R}. \tag{1.3}$$

Theorem 1.1 shows that if one wishes to approximate Euclidean distances in T , it suffices to use a number of hyperplanes that is proportional to the squared Gaussian mean-width of T . And, as was mentioned previously, the Gaussian mean-width is a natural measure of the ‘intrinsic dimension’ of the set. For instance:

- Let E be a d -dimensional subspace and $T = E \cap B_2^n$; then $\ell_*^2(T) \simeq d$.
- Let $T = \Sigma_{s,n}$ be the set of all s -sparse vectors in the Euclidean unit ball. It is standard to verify that $\ell_*^2(T) \simeq \log \binom{n}{s} \simeq s \log(en/s)$. Let B_1^n be the unit ball in ℓ_1^n and recall that $\text{conv}(\Sigma_{s,n})$ is equivalent to $\sqrt{s} B_1^n \cap B_2^n$ (see [20, Lemma 3.1]), the set of approximately s -sparse vectors in the Euclidean unit ball. Thus, Theorem 1.1 implies that only

$$c(L) \frac{\log(2/\rho)}{\rho^3} s \log(en/s)$$

random hyperplanes are needed to approximate distances in $\sqrt{s} B_1^n \cap B_2^n$.

¹ Recall that a random vector is *isotropic* if it is centred and its covariance matrix is the identity; thus, for every $x \in \mathbb{R}^n$, $\mathbb{E}\langle X, x \rangle^2 = \|x\|_2^2$. A centred random vector is *L -subgaussian* if for every $x \in \mathbb{R}^n$ and $p \geq 2$, $\|\langle X, x \rangle\|_{L^p} \leq L\sqrt{p} \|\langle X, x \rangle\|_{L^2}$. Thus, the ψ_2 norms and the L^2 norms of linear forms are equivalent.

Remark 1.2. Note that the lower estimate in (1.3) implies that the hyperplanes endow a ρ -uniform tessellation: any cell of the tessellation of T has diameter at most ρ .

In the second scenario the focus is on *heavy-tailed* random vectors: again X is isotropic and symmetric, but in addition one only assumes that linear forms satisfy an L^1 - L^2 equivalence:

$$\| \langle X, x \rangle \|_{L^2} \leq L \| \langle X, x \rangle \|_{L^1} \quad \text{for every } x \in \mathbb{R}^n. \tag{1.4}$$

In the heavy-tailed scenario a different complexity parameter dictates the required number of hyperplanes. Let X_1, \dots, X_m be independent copies of X and for $K \subset \mathbb{R}^n$ set

$$E(K) := \mathbb{E} \sup_{x \in K} \left| \left\langle \frac{1}{\sqrt{m}} \sum_{i=1}^m \varepsilon_i X_i, x \right\rangle \right|,$$

where $(\varepsilon_i)_{i \geq 1}$ is a sequence of independent, symmetric $\{-1, 1\}$ -valued random variables that is independent of X_1, \dots, X_m .

Remark 1.3. If X_1, \dots, X_m happen to be isotropic, symmetric and L -subgaussian, then $E(K) \leq c(L)\ell_*(K)$ for a constant c that depends only on L . This is one of the features of subgaussian processes and an outcome of Talagrand’s majorizing measures theorem [25]. However, finding upper bounds on $E(K)$ when X is not subgaussian is a challenging question that has been studied extensively over the last 30 years or so and which will not be pursued here.

Theorem 1.4 is a special case of Theorem 2.2 below. In what follows, given $K \subset \mathbb{R}^n$ and $r > 0$, denote by $\mathcal{N}(K, r)$ the smallest number of Euclidean balls of radius r that are needed to cover K .

Theorem 1.4. *There exist constants c_0, \dots, c_4 that depend only on L for which the following holds. Fix $0 < \rho < R$, let $T \subset RB_2^n$ and set $U = \text{conv}(T)$. Let $\lambda = c_0 R$, $r = c_1 \rho^2 / R$, $U_r = (U - U) \cap rB_2^n$ and assume that*

$$m \geq c_2 \left(\left(\frac{R E(U_r)}{\rho^2} \right)^2 + \frac{R \log \mathcal{N}(U, r)}{\rho} \right).$$

Then with probability at least $1 - 2 \exp(-c_3 m (\rho/R)^2)$, for every $x, y \in U$ that satisfy $\|x - y\|_2 \geq \rho$,

$$c_3 \frac{\|x - y\|_2}{R} \leq d(x, y) \leq c_4 \frac{R}{\rho} \cdot \frac{\|x - y\|_2}{R}. \tag{1.5}$$

Remark 1.5. The upper bound in (1.5) features the factor R/ρ ; it replaces $\sqrt{\log(eR/\rho)}$ which appears in the upper estimate in (1.3). This should come as no surprise: the uniform upper estimate on $d(x, y)$ deteriorates the more ‘heavy-tailed’ the random vector X is. At the same time, the lower bound is universal—reflecting the fact that such lower bounds are due to a small-ball property and have nothing to do with tail estimates.

The universal lower bound implies that almost regardless of the choice of X , if x and y are reasonably ‘far apart’ then their distance is exhibited by the fraction of tessellation hyperplanes that separate the points.

The connection between the number of hyperplanes m and the accuracy ρ is less explicit in Theorem 1.4, because $E(U_r)$ depends on m . And even though the uniform central limit theorem shows that $E(U_r)$ converges to $\ell_*(U_r)$ as m tends to infinity, one is interested in quantitative estimates, which are, in general, nontrivial. Since estimating $E(U_r)$ is not the main focus of this article, we shall not pursue the question of controlling $E(U_r)$ for general sets U any further. Instead, and just to illustrate the outcome of Theorem 1.4, let us consider the set $T = \Sigma_{s,n}$.

Example 1.6. Let $T = \Sigma_{s,n}$ and observe that $U, U_r \subset 4(\sqrt{s} B_1^n \cap B_2^n) \subset 8 \operatorname{conv}(\Sigma_{s,n})$. By Sudakov’s inequality (see, e.g., [16]),

$$\log \mathcal{N}(T, r) \leq c_1 \frac{\ell_*^2(U_r)}{r^2} \leq c_2 \frac{s \log(en/s)}{\rho^4}.$$

Moreover, $E(U_r) \leq 4E(\operatorname{conv}(\Sigma_{s,n})) = 4E(\Sigma_{s,n})$, and there are many generic cases in which

$$E(\Sigma_{s,n}) \lesssim s \log(en/s). \tag{1.6}$$

For example, following [15, 17], one may show that (1.6) holds when X is isotropic, unconditional and log-concave; and also when X has i.i.d. coordinates distributed according to a mean-zero, variance 1 random variable ξ that satisfies $(\mathbb{E}|\xi|^p)^{1/p} \lesssim p^\alpha$ for some $\alpha > 0$ and for every $p \leq \log n$. We refer to [9, Section V] for proofs of these facts and for other examples of a similar nature.

When (1.6) holds, then Theorem 1.4 implies that it is enough to use

$$m = c(L) \frac{s \log(en/s)}{\rho^4}$$

hyperplanes to estimate distances in $\sqrt{s} B_1^n \cap B_2^n$. And although the way m scales with ρ is worse than in the subgaussian case, the scaling with s and n is the same.

Before presenting the proofs of Theorems 1.1 and 1.4, let us explore the connection between random hyperplane tessellations and signal recovery problems. Readers that are solely interested in hyperplane tessellations can safely skip straight to Section 2, where the proofs of the two theorems may be found.

1.1. Application to one-bit compressed sensing

One good reason for studying non-Gaussian random hyperplane tessellations of arbitrary sets comes from signal recovery problems involving quantized measurements. By *quantization* we mean converting analog measurements of a signal into a finite number of bits. This essential step is part of any signal processing procedure and allows one to digitally transmit, process, and reconstruct signals. The area of *quantized compressed sensing* investigates how to design a measurement procedure, quantizer, and reconstruction algorithm that together recover low-complexity signals—such as signals that have a sparse representation in a given basis. An efficient system has to be able to reconstruct signals

based on a minimal number of measurements, each of which is quantized to the smallest number of bits, and to do so via a computationally efficient reconstruction algorithm. In addition, the system should be reliable: it should be robust to both pre-quantization noise (noise in the analog measurements process) and post-quantization noise (bit corruptions that occur during the quantization process).

Our interest here is in the popular *one-bit compressed sensing model*, in which one observes quantized measurements of the form

$$q = \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}), \quad (1.7)$$

where $A \in \mathbb{R}^{m \times n}$, $m \ll n$, sign is the sign function applied elementwise, $v_{\text{noise}} \in \mathbb{R}^m$ is a vector modelling the noise in the analog measurement process and $\tau_{\text{thres}} \in \mathbb{R}^m$ is a (possibly random) vector consisting of quantization thresholds. We restrict ourselves to memoryless quantization, meaning that the thresholds are set in a non-adaptive manner. In this case, the one-bit quantizer $\text{sign}(\cdot + \tau_{\text{thres}})$ can be implemented efficiently in practice, and because of its efficiency it has been very popular in engineering literature—especially in applications in which analog-to-digital converters represent a significant factor in the energy consumption of the measurement system (see e.g. [5, 18]).

In spite of its popularity, there are only a few rigorous results that show that one-bit compressed sensing is viable: the vast majority of mathematical literature (see e.g. [3, 13, 14, 20, 21]) on one-bit compressed sensing has focused on the special case in which A is a standard Gaussian matrix, and the practical relevance of such results is limited—Gaussian matrices cannot be realized in a real-world measurement setup. As an additional difficulty, it is well known that one-bit compressed sensing may perform poorly outside the Gaussian setup. In fact, it can very easily *fail*, even if the measurement matrix is known to perform optimally in ‘unquantized’ compressed sensing. For example, if the threshold vector τ_{thres} is zero, there are 2-sparse vectors that cannot be distinguished based on their one-bit Bernoulli measurements (see Figure 3).

As an application of the new hyperplane tessellation results described in the previous section, we show that one-bit compressed sensing can actually perform well in scenarios that are far more general than the Gaussian setting. What makes all the difference is the rather striking effect that dithering (that is, adding well-designed ‘noise’ to the measurements before quantizing) has on the one-bit quantizer. Indeed, thanks to dithering, accurate recovery from one-bit measurements is possible even if the measurement vectors are drawn from a heavy-tailed distribution. Moreover, the recovery results are robust to both adversarial and potentially heavy-tailed stochastic noise on the analog measurements, as well as to adversarial bit corruptions that may occur during quantization.

In what follows we explain why dithering has such an effect: the geometric interpretation of dithering leads to random tessellations that can be used to approximate distances between signals. The ability to approximate distances has a crucial impact on the performance of recovery procedures.

To understand the connection between hyperplane tessellations and signal recovery from one-bit quantized measurements, let us first assume that no bit corruptions occur in the quantization process, and that there is no pre-quantization noise ($v_{\text{noise}} = 0$). In this case, one observes $q = \text{sign}(Ax + \tau_{\text{thres}})$. If X_1, \dots, X_m denote the rows of A and




Fig. 3. Symmetric Bernoulli vectors in \mathbb{R}^2 can only generate two different homogeneous hyperplanes. As a result, there exist two points on the sphere (here, e_1 and $(e_1 + \lambda e_2)/\sqrt{1 + \lambda^2}$ for $-1 < \lambda < 0$, both marked in red) that are far apart, but cannot be separated by a Bernoulli hyperplane. This problem persists in high dimensions. In addition, any two points lying on a ray originating from 0 (e.g., the points that are marked in green) cannot be separated by a homogeneous hyperplane (the latter problem is not specific to the Bernoulli case). Both problems can be solved by using parallel shifts of the hyperplanes instead of the homogeneous ones.

τ_1, \dots, τ_m are the entries of τ_{thres} , then q encodes the cell of the hyperplane tessellation in which the signal x is located. A popular strategy used for recovering x is searching for a vector $x^\# \in T$ that is *quantization consistent*, i.e., $q = \text{sign}(Ax^\# + \tau_{\text{thres}})$. For instance, if $T = \Sigma_{s,n}$, the set of all s -sparse vectors in the Euclidean unit ball, then one can find such a vector by solving

$$\min_{z \in \mathbb{R}^n} \|z\|_0 \quad \text{s.t.} \quad q = \text{sign}(Az + \tau_{\text{thres}}), \quad \|z\|_2 \leq 1. \tag{1.8}$$

Geometrically, a quantization consistent vector is simply a vector lying in the same cell as x , and one can ensure that $\|x^\# - x\|_2 \leq \rho$ by showing that $\|x - y\|_2 \leq \rho$ for any $y \in T$ located in the same cell as x . Since there is no additional information on the identity of the cell in which x is located, one has to ensure that any pair of points in T located in the same cell are at distance at most ρ from each other, i.e., the hyperplanes H_{X_i, τ_i} must form a ρ -uniform tessellation of T . Phrased differently, if $x, y \in T$ are at distance at least ρ , then that fact must be exhibited by the hyperplanes H_{X_i, τ_i} : at least one of the hyperplanes must separate x and y . In particular, if one has access to a ρ -uniform tessellation of T , one can uniformly recover signals from T using only $\text{sign}(Ax + \tau_{\text{thres}})$ as data. Moreover, the reverse direction is clearly true: the degree of accuracy in uniform recovery results in T is determined by the largest diameter (in T) of a cell of the tessellation formed by the hyperplanes H_{X_i, τ_i} .

Unfortunately, even if $(H_{X_i, \tau_i})_{i=1}^m$ forms a uniform tessellation of T there is still the question of pre- and post-quantization noise one has to contend with. To understand the effect of post-quantization noise (i.e., bit corruptions that occur during quantization), assume that one observes a corrupted sequence of bits $q_{\text{corr}} \in \{-1, 1\}^m$, where the i -th bit being corrupted means that instead of receiving $q_i = \text{sign}(\langle X_i, x \rangle + \tau_i)$ from the quan-

tizer, one observes $(q_{\text{corr}})_i = -\text{sign}(\langle X_i, x \rangle + \tau_i)$; thus, one is led to believe that x is on the ‘wrong side’ of the i -th hyperplane H_{X_i, τ_i} . As a consequence, recovery methods that search for a quantization consistent vector can easily fail even if a single bit is corrupted. For instance, the program (1.8) (with q replaced by q_{corr}) will, in the best case scenario, search for a vector in the wrong cell of the tessellation, and in the worse case, the corrupted bit may cause a conflict and there will be no sparse vector z satisfying $q_{\text{corr}} = \text{sign}(Az + \tau_{\text{thres}})$ (see Figure 4 for an illustration).




Fig. 4. The effect of a bit corruption associated with the dashed, red hyperplane H_{X_i, τ_i} . Either the bit corruption leads the program (1.8) (with q replaced by q_{corr}) to search in the wrong cell of the tessellation marked by the red dot (left) or causes the program to be infeasible (right).

The effect of pre-quantization noise (i.e., noise in the analog measurement process) is equally problematic: noise simply causes a parallel shift of the hyperplane H_{X_i, τ_i} , and one has no control over the size of this ‘noise-induced’ shift. Again, the recovery program (1.8) (with $q = \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}})$) can easily fail if pre-quantization noise is present (see Figure 5).




Fig. 5. The effect of a noise-induced parallel shift of the dashed, blue hyperplane H_{X_i, τ_i} onto the dashed, red hyperplane $H_{X_i, v_i + \tau_i}$. The program (1.8) (with $q = \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}})$) searches for a vector z with $\text{sign}(\langle X_i, z \rangle + \tau_i) = \text{sign}(\langle X_i, z \rangle + v_i + \tau_i)$. This means that the program incorrectly searches for a solution located to the right of the dashed, blue hyperplane H_{X_i, τ_i} ; as a consequence, a solution is found in the wrong cell of the tessellation marked by the red dot (left) or it can even happen that no feasible point exists (right).

One possible way of overcoming this ‘infeasibility problem’ due to noise is by designing a recovery program that is stable: its output does not change by much even if some of the given bits are misleading. For example, one may try searching for a vector $z \in T$ whose uncorrupted quantized measurements $\text{sign}(Az + v_{\text{noise}} + \tau_{\text{thres}})$ are closest to the observed corrupted vector q_{corr} . However, since one does not have access to v_{noise} , one can only try to match its proxy $\text{sign}(Az + \tau_{\text{thres}})$ to q_{corr} , i.e., to solve

$$\min_{z \in \mathbb{R}^n} d_H(q_{\text{corr}}, \text{sign}(Az + \tau_{\text{thres}})) \quad \text{s.t.} \quad z \in T, \tag{1.9}$$

where d_H denotes the Hamming distance. In the context of sparse recovery, the latter program is

$$\min_{z \in \mathbb{R}^n} d_H(q_{\text{corr}}, \text{sign}(Az + \tau_{\text{thres}})) \quad \text{s.t.} \quad \|z\|_0 \leq s, \quad \|z\|_2 \leq 1. \tag{1.10}$$

Remark 1.7. Note that this program requires (a good estimate of) the signal sparsity as input, in contrast to (1.8).

To ensure that (1.9) yields an accurate reconstruction, the uniform tessellation has to be finer than in the corruption-free case: even if some signs are ‘flipped’, the distance between points in the resulting cell and points in the true one should still be small. And indeed, our results ensure that the hyperplane tessellation is sufficiently fine: for any $x, y \in T$ that are at least ρ -separated there are *many* hyperplanes that separate the two points—of the order of $\|x - y\|_2 m$. Thus, even after corrupting $\simeq \rho m$ bits one may still detect that x and y are ‘far away’ from one another.

Finally, although (1.9) can guarantee robust signal recovery, there are no guarantees that it can be solved efficiently. In addition, since (1.9) matches $\text{sign}(Az + \tau_{\text{thres}})$, rather than $\text{sign}(Az + v_{\text{noise}} + \tau_{\text{thres}})$, to q_{corr} , it is still quite sensitive to pre-quantization noise. Both problems can be mended by convexification. Indeed, observe that

$$d_H(q_{\text{corr}}, \text{sign}(Az + v_{\text{noise}} + \tau_{\text{thres}})) = \frac{1}{2} \sum_{i=1}^m (1 - (q_{\text{corr}})_i \text{sign}(\langle X_i, z \rangle + v_i + \tau_i)).$$

One may relax this objective function by replacing $\text{sign}(\langle X_i, z \rangle + v_i + \tau_i)$ by $\langle X_i, z \rangle + v_i + \tau_i$ and relax the constraint $z \in T$ to $z \in \text{conv}(T)$ leading to the convex program

$$\min_{z \in \mathbb{R}^n} \frac{1}{2} \sum_{i=1}^m (1 - (q_{\text{corr}})_i (\langle X_i, z \rangle + v_i + \tau_i)) \quad \text{s.t.} \quad z \in \text{conv}(T).$$

An equivalent formulation of this program, which only requires the known data q_{corr} and A , is

$$\max_{z \in \mathbb{R}^n} \frac{1}{m} \langle q_{\text{corr}}, Az \rangle \quad \text{s.t.} \quad z \in \text{conv}(T), \tag{1.11}$$

and in contrast to (1.9), (1.11) does not require the threshold vector τ_{thres} as input.

The recovery program (1.11) was proposed in [21]; and in what follows we explore a regularized version of that program: for $\lambda > 0$ consider

$$\max_{z \in \mathbb{R}^n} \frac{1}{m} \langle q_{\text{corr}}, Az \rangle - \frac{1}{2\lambda} \|z\|_2^2 \quad \text{s.t.} \quad z \in \text{conv}(T), \tag{1.12}$$

which, in the context of sparse recovery, corresponds to the tractable program

$$\max_{z \in \mathbb{R}^n} \frac{1}{m} \langle q_{\text{corr}}, Az \rangle - \frac{1}{2\lambda} \|z\|_2^2 \quad \text{s.t.} \quad \|z\|_1 \leq \sqrt{s}, \|z\|_2 \leq 1.$$

Remark 1.8. We refer the reader to [21] for an extensive discussion of the connections between the recovery program (1.11) and the literature on regression with a binary response variable.

Let us formulate the main signal recovery results of this article, which are direct outcomes of the results on random tessellations.

Fix a target reconstruction error ρ , recall that the quantization thresholds τ_i are i.i.d. uniformly distributed in $[-\lambda, \lambda]$, assume that the entries v_i of v_{noise} are i.i.d. copies of a random variable v and that at most βm of the bits are arbitrarily corrupted during quantization, i.e., the observed corrupted vector q_{corr} satisfies $d_H(q_{\text{corr}}, q) \leq \beta m$. The adversarial component of the pre-quantization noise v is $|\mathbb{E}v|$, σ^2 is its variance and $\|v\|_{L^2}$ is its L^2 norm. We write $T_r = (T - T) \cap rB_2^n$ for any $r > 0$.

The first recovery result concerns the recovery program (1.9) in the L -subgaussian scenario, in which the rows X_i of A are i.i.d. copies of a symmetric, isotropic, L -subgaussian vector X . In addition, assume that v also satisfies $\|v\|_{L^p} \leq L\sqrt{p} \|v\|_{L^2}$ for every $p \geq 2$.

Theorem 1.9. *There exist constants $c_0, \dots, c_4 > 0$ depending only on L such that the following holds. Let $T \subset RB_2^n$, set $\lambda \geq c_0(R + \|v\|_{L^2}) + \rho$ and put $r = c_1\rho/\sqrt{\log(e\lambda/\rho)}$. Assume that*

$$m \geq c_2\lambda \left(\frac{\ell_*^2(T_r)}{\rho^3} + \frac{\log \mathcal{N}(T, r)}{\rho} \right),$$

and that $|\mathbb{E}v| \leq c_3\rho$, $\sigma \leq c_3\rho/\sqrt{\log(e\lambda/\rho)}$ and $\beta \leq c_3\rho/\lambda$. Then with probability at least $1 - 2 \exp(-c_4m\rho/\lambda)$, for every $x \in T$, any solution $x^\#$ of (1.9) satisfies $\|x^\# - x\|_2 \leq \rho$.

Example 1.10. To put Theorem 1.9 in some context, consider an arbitrary $T \subset B_2^n$ and assume $\|v\|_{L^2} \leq 1$, so that λ is a constant that depends only on L . By Sudakov’s inequality,

$$\log \mathcal{N}(T, r) \leq c \frac{\ell_*^2(T)}{r^2} \leq c(L) \frac{\log(e/\rho)}{\rho^2} \ell_*^2(T), \tag{1.13}$$

and trivially $\ell_*(T_r) \leq \ell_*(T)$, which means that a sample size of

$$m = c'(L) \frac{\log(e/\rho)}{\rho^3} \ell_*^2(T)$$

suffices for recovery. In the special case of $T = \Sigma_{s,n}$ a much better estimate is possible. Indeed, it is standard to verify that there is an absolute constant c such that for any $1 \leq s \leq n$,

$$\ell_*(\Sigma_{s,n}) \simeq \sqrt{s \log(en/s)} \quad \text{and} \quad \log \mathcal{N}(\Sigma_{s,n}, r) \leq cs \log\left(\frac{en}{sr}\right). \tag{1.14}$$

Moreover, since $(\Sigma_{s,n} - \Sigma_{s,n}) \cap rB_2^n \subset r\Sigma_{2s,n}$ it follows that

$$\ell_*(T_r) \leq cr\sqrt{s \log(en/s)} = c(L) \frac{\rho}{\sqrt{\log(e/\rho)}} \cdot \sqrt{s \log(en/s)},$$

implying that a sample size of

$$m = c'(L)\rho^{-1}s \log\left(\frac{en}{s\rho}\right) \tag{1.15}$$

guarantees that with high probability one can recover any s -sparse vector in B_2^n with accuracy ρ via (1.9).

In the heavy-tailed scenario, one only assumes that X is isotropic, symmetric, and satisfies the L^1 - L^2 equivalence (1.4), and that v has finite variance σ^2 and satisfies an L^1 - L^2 equivalence.

Theorem 1.11. *There exist constants $c_0, \dots, c_4 > 0$ depending only on L such that the following holds. Assume that $T \subset RB_2^n$. Let $\lambda \geq c_0(R + \|v\|_{L^2}) + \rho$, set $r = c_1\rho^2/\lambda$, and suppose that*

$$m \geq c_2 \left(\left(\frac{\lambda E(T_r)}{\rho^2} \right)^2 + \lambda \frac{\log \mathcal{N}(T, r)}{\rho} \right). \tag{1.16}$$

Assume further that $|\mathbb{E}v| \leq c_3\rho$, $\sigma \leq c_3\rho^{3/2}/\sqrt{\lambda}$ and $\beta \leq c_3\rho/\lambda$. Then with probability at least $1 - 2 \exp(-c_4m(\rho/\lambda)^2)$, for every $x \in T$, any solution $x^\#$ of (1.9) satisfies $\|x^\# - x\|_2 \leq \rho$.

Example 1.12. To illustrate the outcome of Theorem 1.11, assume that $\|v\|_{L^2} \leq 1$ and consider $T = \Sigma_{s,n}$; hence, λ is a constant that depends only on L . Since $T_r \subset r\Sigma_{2s,n}$, the first term in (1.16) is bounded by $E^2(\Sigma_{2s,n})$. As noted previously, there are many natural random vectors that are more heavy-tailed than subgaussian, and still $E^2(\Sigma_{2s,n}) \simeq s \log(en/s)$. In such cases, the sample size (1.15) is sufficient for recovery.

Let us compare Theorems 1.9 and 1.11 to existing work. As was mentioned previously, almost all the signal reconstruction results in (memoryless) one-bit compressed sensing are based on the assumption that the measurement matrix is Gaussian (see e.g. [8] for an overview). Among those, the work that is closest to ours is [13], where there is no dithering involved in the recovery procedure ($\tau_{\text{thres}} = 0$) and thus it is only possible to recover signals located on the unit sphere. It was shown in [13, Theorem 2] that if $A \in \mathbb{R}^{m \times n}$ is standard Gaussian and $m \gtrsim \rho^{-1}s \log(n/\rho)$ then, with high probability,

any s -sparse $x, x' \in S^{n-1}$ for which $\text{sign}(Ax) = \text{sign}(Ax')$ satisfy $\|x - x'\|_2 \leq \rho$. In particular, one can approximate x with accuracy ρ by solving the nonconvex program

$$\min_{z \in \mathbb{R}^n} \|z\|_0 \quad \text{s.t.} \quad \text{sign}(Ax) = \text{sign}(Az), \|z\|_2 = 1.$$

In comparison, Theorem 1.9 shows that a similar result holds in the subgaussian scenario—and at the same time extends it to sparse vectors in the unit ball and makes it robust to pre- and post-quantization noise. Clearly, such a generalization is possible thanks to the effect of dithering. Remarkably, Theorem 1.11 shows that this result can be extended further to a large class of heavy-tailed measurements. In fact, Theorem 1.11 is the first recovery result of its kind—involving quantized measurements that can be heavy-tailed.

In [3, 14] the authors study sparse recovery with Gaussian measurements and introduce standard Gaussian dithering to derive recovery results for sparse vectors in the Euclidean unit ball. The idea behind these results is to use a ‘lifting trick’: for instance, in [3] one interprets the dithered measurements $\text{sign}(Ax + \tau)$ as $\text{sign}([A \ \tau][x, 1]/\|[x, 1]\|_2)$, where $[A \ \tau]$ is obtained by appending τ to A as an additional column. Since $[A \ \tau]$ is a standard Gaussian again, recovery methods for sparse vectors on the Euclidean unit sphere can be used to find an approximation of $[x, 1]/\|[x, 1]\|_2$ of the form $[x^\#, 1]/\|[x^\#, 1]\|_2$. Afterwards, one can bound $\|x - x^\#\|_2$ by the distance between the last two vectors. Since this lifting argument is based on a reduction to the one-bit compressed sensing with zero thresholds model, it ‘imports’ the strong limitations of that model; in particular, it cannot be used to derive recovery results for non-Gaussian measurements. In addition, since the recovery methods in [3, 14] rely on enforcing quantization consistency, they are not robust to post-quantization noise. In contrast, thanks to the geometric interpretation of dithering, the recovery results presented here are robust, hold for non-Gaussian measurements matrices and for general signal sets.

Finally, let us formulate the main recovery result for the program (1.12) in the L -subgaussian scenario. Here, v is centred and L -subgaussian with variance σ^2 . Set $U = \text{conv}(T)$ and $U_\rho = (U - U) \cap \rho B_2^n$.

Theorem 1.13. *There exist constants c_0, \dots, c_4 that depend only on L for which the following holds. Let $T \subset RB_2^n$, fix $\rho > 0$, set*

$$\lambda \geq c_0(\sigma + R)\sqrt{\log(c_0(\sigma + R)/\rho)}$$

and let $r = c_1\rho/\log(e\lambda/\rho)$. If m and β satisfy

$$m \geq c_2 \left(\left(\frac{\lambda \ell_*(U_\rho)}{\rho^2} \right)^2 + \lambda^2 \frac{\log \mathcal{N}(T, r)}{\rho^2} \right), \quad \beta \sqrt{\log(e/\beta)} = c_3 \frac{\rho}{\lambda},$$

then, with probability at least $1 - 2 \exp(-c_4 m \rho^2 / \lambda^2)$, for any $x \in T$ the solution $x^\#$ of (1.12) satisfies $\|x^\# - x\|_2 \leq \rho$.

Example 1.14. Let $T = \sqrt{s} B_1^n \cap B_2^n$ and assume that $\sigma \leq 1$. Observe that $T = U$ and that one may set $\lambda = c_0(L)\sqrt{\log(e/\rho)}$. Also, for $0 < \rho \leq 1$, $U_\rho \subset 2(\sqrt{s} B_1^n \cap \rho B_2^n)$, and it is standard to verify that $\ell_*(U_\rho) \simeq \sqrt{s} \max\{\log(en\rho^2/s), 1\}$. Taking the estimate (1.13) for $\log \mathcal{N}(T, r)$ into account, it is evident that if

$$m = c(L) \frac{s \log(en/s) \log^3(e/\rho)}{\rho^4}$$

then with high probability one may recover any $x \in T$ using the convex recovery procedure (1.12), even in the presence of pre- and post-quantization noise.

In the context of Gaussian measurement matrices, Theorem 1.13 improves upon the work of Plan and Vershynin [21], who considered the situation when there is no dithering ($\tau_{\text{thres}} = 0$). They introduced the convex program (1.11) and proved recovery results for signal sets $T \subset S^{n-1}$ of two different flavours. In a nonuniform recovery setting² they showed that $m \gtrsim \rho^{-4} \ell_*^2(T)$ measurements suffice to reconstruct a fixed signal, even if pre-quantization noise is present and quantization bits are randomly flipped with a probability that is allowed to be arbitrarily close to 1/2. In the uniform recovery setting, they showed that if $m \gtrsim \rho^{-12} \ell_*^2(T)$, one can achieve a reconstruction error ρ even if a fraction $\beta = \rho^2$ of the received bits are corrupted in an adversarial manner while quantizing. Theorem 1.13 extends the latter result to subgaussian measurements with a better condition on m and β , and at the same time incorporates pre-quantization noise and allows the reconstruction of signals that need not be located on the unit sphere.

As noted previously, there are very few reconstruction results available when the measurements are not standard Gaussian. The work [1] generalizes the nonuniform recovery results from [21] to subgaussian measurements under additional restrictions. For $T \subset S^{n-1}$ and a fixed $x \in T$ it is shown that $m \gtrsim \rho^{-4} \ell_*^2(T)$ measurements suffice to reconstruct x up to error ρ via (1.11), provided that either $\|x\|_\infty \leq \rho^4$ (meaning that the signal must be well-spread) or the total variation distance between the subgaussian measurements and the standard Gaussian distribution is at most ρ^{16} . Theorem 1.13 is a considerable improvement of those results.

Remark 1.15. At the expense of substantial additional technicalities, the proof strategies developed in this work lead to recovery results for sparse vectors when A is a random partial circulant matrix generated by a subgaussian random vector. The latter model occurs in several practical measurement setups, including SAR radar imaging, Fourier optical imaging and channel estimation (see e.g. [24] and the references therein). To keep this work accessible to a general audience and in an attempt to clearly present the main ideas used in the proofs, we choose to defer the additional technical developments needed for the circulant case to a companion work [10].

² In the uniform recovery setting one attains a high probability event on which recovery is possible for all $x \in T$, whereas in nonuniform recovery the event depends on the signal $x \in T$.

1.2. Notation

We use $\|x\|_p$ to denote the ℓ_p norm of $x \in \mathbb{R}^n$ and B_p^n denotes the ℓ_p -unit ball in \mathbb{R}^n . For a subgaussian random variable ξ let

$$\|\xi\|_{\psi_2} := \sup_{p \geq 1} \frac{\|\xi\|_{L^p}}{\sqrt{p}};$$

thus, a random variable is L -subgaussian in the sense stated in the introduction precisely when $\|\xi\|_{\psi_2} \leq L\|\xi\|_{L^2}$. This is equivalent to

$$\mathbb{P}(|\xi| \geq t) \leq c_1 e^{-c_2 t^2 / (L\|\xi\|_{L^2})^2}, \quad t \geq 0,$$

for some absolute constants $c_1, c_2 > 0$.

In what follows, \mathcal{U} denotes the uniform distribution. For $k \in \mathbb{N}$ set $[k] = \{1, \dots, k\}$ and for a set S let $|S|$ denote its cardinality. d_H is the (unnormalized) Hamming distance on the discrete cube and $\Sigma_{s,n} = \{x \in \mathbb{R}^n : \|x\|_0 \leq s, \|x\|_2 \leq 1\}$ is the set of s -sparse vectors in the Euclidean unit ball. For $T \subset \mathbb{R}^n$ set $T_r = (T - T) \cap rB_2^n$ and denote by $\text{conv}(T)$ its convex hull. The Gaussian mean-width of T is denoted by $\ell_*(T)$ and for any $r > 0$ let $\mathcal{N}(T, r)$ be the smallest number of Euclidean balls of radius r that are needed to cover T . Finally, c and C denote absolute constants; their value may change from line to line. c_α or $c(\alpha)$ are constants that depend only on the parameter α , $a \lesssim_\alpha b$ implies that $a \leq c_\alpha b$, and $a \simeq_\alpha b$ means that both $a \lesssim_\alpha b$ and $a \gtrsim_\alpha b$ hold.

2. Random tessellations

This section is devoted to the proof of our main tessellation results, Theorems 2.2 and 2.3, which are generalizations of Theorems 1.1 and 1.4 respectively.

Before formulating the results let us define a mild structural property of a subset of a metric space.

Definition 2.1. Let (\mathcal{X}, d) be a metric space. A set $T \subset \mathcal{X}$ is (r, γ) -metrically convex in \mathcal{X} if for every $x, y \in T$ with $d(x, y) \geq r$ there are $z_1, \dots, z_\ell \in \mathcal{X}$ such that

$$\gamma r \leq d(z_i, z_{i+1}) \leq r \quad \text{and} \quad \sum_{i=0}^{\ell} d(z_i, z_{i+1}) \leq \gamma^{-1} d(x, y),$$

where we set $z_0 = x, z_{\ell+1} = y$. If $\mathcal{X} = T$, then we say that T is (r, γ) -metrically convex.

The idea behind this notion is straightforward: it implies that controlling ‘local oscillations’ of a function f ensures that it satisfies a Lipschitz condition for long distances. Indeed, assume that $\sup_{\{w, v \in \mathcal{X} : d(w, v) \leq r\}} |f(w) - f(v)| \leq \kappa$ and for any $x, y \in T$ that satisfy $d(x, y) \geq 2r$ let $(z_i)_{i=0}^{\ell+1}$ be as in Definition 2.1. Then

$$\begin{aligned} |f(x) - f(y)| &\leq \left| \sum_{i=0}^{\ell} (f(z_i) - f(z_{i+1})) \right| \leq \kappa(\ell + 1) \leq \frac{\kappa}{\gamma r} \sum_{i=0}^{\ell} d(z_i, z_{i+1}) \\ &\leq \frac{\kappa}{\gamma^2 r} d(x, y). \end{aligned} \tag{2.1}$$

Therefore, f satisfies a Lipschitz condition for long distances with constant $\kappa/(\gamma^2 r)$.

Observe that if T is a convex subset of a normed space then it is $(r, 1)$ -metrically convex for any $r > 0$; also, every subset of a normed space is $(r, 1)$ -metrically convex in its convex hull. Finally, $\Sigma_{s,n}$ is (r, γ) -metrically convex in $\Sigma_{2s,n}$ for an absolute constant γ . We omit the straightforward proofs of these claims.

Let us first state the main result in the heavy-tailed scenario. Consider a random vector X that is isotropic, symmetric, and satisfies an L^1 - L^2 norm equivalence: for every $t \in \mathbb{R}^n$,

$$\|t\|_2 = \|\langle X, t \rangle\|_{L^2} \leq L \|\langle X, t \rangle\|_{L^1}. \tag{2.2}$$

Theorem 2.2. *There exist constants c_0, \dots, c_4 that depend only on L for which the following holds. Let $T \subset RB_2^n$ and set $\lambda \geq c_0R$. Suppose that $0 < r < \rho < \lambda$ satisfy $r \leq c_1\rho^2/\lambda$ and assume that*

$$\log \mathcal{N}(T, r) \leq c_2 \frac{m\rho}{\lambda} \quad \text{and} \quad E(T_r) \leq c_2 \frac{\rho^2}{\lambda} \sqrt{m}.$$

Then with probability at least $1 - 8 \exp(-c_3m(\rho/\lambda)^2)$, for every $x, y \in T$ that satisfy $\|x - y\|_2 \geq \rho$,

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \geq c_4m \frac{\|x - y\|_2}{\lambda}.$$

Moreover, if T is (r, γ) -metrically convex then on the same event, if $\|x - y\|_2 \geq 2r$,

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \leq \frac{c_5\lambda}{\rho\gamma^2} \cdot m \frac{\|x - y\|_2}{\lambda}.$$

Proof of Theorem 1.4. Apply Theorem 2.2 to the set $U = \text{conv}(T)$, which is $(r, 1)$ -metrically convex for any $r > 0$, and for the parameters $\lambda = c_0R$ and $r = c_1\rho^2/R$. With these choices Theorem 1.4 follows immediately. \square

When X is L -subgaussian one may establish a sharper result.

Theorem 2.3. *There exist constants c_0, \dots, c_5 that depend only on L for which the following holds. Let $T \subset RB_2^n$, set $\lambda \geq c_0R$ and consider an isotropic, symmetric, L -subgaussian random vector X . Let m and $0 < r < \rho < \lambda$ satisfy*

$$\rho \geq c_1r\sqrt{\log(e\lambda/\rho)}, \quad m \geq c_2 \max \left\{ \frac{\lambda}{\rho} \log \mathcal{N}(T, r), \lambda \frac{\ell_*^2(T_r)}{\rho^3} \right\}.$$

Then with probability at least $1 - 8 \exp(-c_3m\rho/\lambda)$, for all $x, y \in T$ such that $\|x - y\|_2 \geq \rho$, one has

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \geq c_4m \frac{\|x - y\|_2}{\lambda}.$$

Moreover, if T is (r, γ) -metrically convex then on the same event, if $\|x - y\|_2 \geq 2r$,

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \leq \frac{c_5\sqrt{\log(e\lambda/\rho)}}{\gamma^2} \cdot m \frac{\|x - y\|_2}{\lambda}.$$

Proof of Theorem 1.1. Theorem 1.1 is an immediate outcome of Theorem 2.3 for $U = \text{conv}(T)$. Indeed, $\text{conv}(T)$ is $(r, 1)$ -metrically convex for any $r > 0$, $\ell_*(U_r) \leq \ell_*(T)$, and by Sudakov’s inequality, $\log \mathcal{N}(U, r) \leq c\ell_*^2(T)/r^2$. The claim follows by setting $r = c\rho/\sqrt{\log(e\lambda/\rho)}$ and $\lambda = c'R$ for suitable absolute constants c and c' . \square

In the context of tessellations, Theorem 2.2 and the first part of Theorem 2.3 improve the estimate from (1.2) in several ways: firstly, Theorem 2.2 holds for a very general collection of random vectors: X has to satisfy a small-ball condition rather than being Gaussian. Secondly, both are valid for any subset of \mathbb{R}^n and not just for subsets of the sphere; and, finally, if X happens to be L -subgaussian, it yields the best known estimate on the diameter of each ‘cell’ in the random tessellation—even when X is Gaussian and T is a subset of S^{n-1} .

2.1. The heavy-tailed scenario

A fundamental question that is at the heart of our arguments has to do with stability: given two points x and y , how ‘stable’ is the set

$$\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\} = (*)$$

to perturbations? If one believes that the cardinality of $(*)$ reflects the distance $\|x - y\|_2$, it stands to reason that if r is significantly smaller than $\|x - y\|_2$ and $\|x - x'\|_2 \leq r$, $\|y - y'\|_2 \leq r$, then $|\{i : \text{sign}(\langle X_i, x' \rangle + \tau_i) \neq \text{sign}(\langle X_i, y' \rangle + \tau_i)\}|$ should not be very different from $|(*)|$.

Unfortunately, stability is not true in general. If either x or y are ‘too close’ to many of the separating hyperplanes, then even a small shift in either one of them can have a dramatic effect on the signs of $\langle X_i, \cdot \rangle + \tau_i$ and destroy the separation. Thus, to ensure stability one requires a stronger property than mere separation: points need to be separated by a large margin.

Definition 2.4. The hyperplane H_{X_i, τ_i} θ -well-separates x and y if

- $\text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)$,
- $|\langle X_i, x \rangle + \tau_i| \geq \theta\|x - y\|_2$, and
- $|\langle X_i, y \rangle + \tau_i| \geq \theta\|x - y\|_2$.

Denote by $I_{x,y}(\theta) \subset [m]$ the set of indices for which H_{X_i, τ_i} θ -well-separates x and y .

The condition that $|\langle X_i, x \rangle + \tau_i|, |\langle X_i, y \rangle + \tau_i| \geq \theta\|x - y\|_2$ is precisely what ensures that perturbations of x or y of the order of $\|x - y\|_2$ do not spoil the fact that the hyperplane H_{X_i, τ_i} separates the two points.

We begin by showing that even in the heavy-tailed scenario and with high probability, $|I_{x,y}(\theta)|$ is proportional to $m\|x - y\|_2$ for any two (fixed) points x and y . Let us stress that the high probability estimate is crucial: it will lead to uniform control on a net of large cardinality.

Theorem 2.5. *There are constants c_1, \dots, c_4 that depend only on L for which the following holds. Let $x, y \in RB_2^n$ and set $\lambda \geq c_1 R$. With probability at least*

$$1 - 4 \exp(-c_2 m \min\{\|x - y\|_2/\lambda, 1\}),$$

we have

$$|I_{x,y}(c_3)| \geq c_4 m \|x - y\|_2/\lambda.$$

The proof of Theorem 2.5 requires two preliminary observations. Consider a random variable τ that satisfies the small-ball estimate

$$\sup_{u \in \mathbb{R}} \mathbb{P}(|\tau - u| \leq \varepsilon) \leq C_\tau \varepsilon \quad \text{for all } \varepsilon \geq 0, \tag{2.3}$$

and let Z be independent of τ . Then clearly

$$\mathbb{P}(|Z + \tau| \leq \varepsilon) \leq C_\tau \varepsilon \quad \text{for all } \varepsilon \geq 0. \tag{2.4}$$

If $\tau \sim \mathcal{U}[-\lambda, \lambda]$ then (2.3) holds for $C_\tau = 1/\lambda$. Therefore, by the Chernoff bound, if $(Z_i)_{i=1}^m$ and $(\tau_i)_{i=1}^m$ are independent copies of Z and τ respectively, then with probability at least $1 - 2 \exp(-cm\varepsilon/\lambda)$,

$$|\{i : |Z_i + \tau_i| \geq \varepsilon\}| \geq \left(1 - \frac{2\varepsilon}{\lambda}\right)m. \tag{2.5}$$

The second observation is somewhat more involved. Consider a random variable τ that satisfies

$$\mathbb{P}(\alpha < \tau \leq \beta) \geq c_\tau(\beta - \alpha) \tag{2.6}$$

for all $-\lambda \leq \alpha \leq \beta \leq \lambda$. Let Z and W be square integrable whose difference satisfies a small-ball condition: there are constants κ and δ such that

$$\mathbb{P}(|Z - W| \geq \kappa \|Z - W\|_{L^1}) \geq \delta.$$

Lemma 2.6. *There are absolute constants c_0 and c_1 and constants $c_2, c_3 \simeq c_\tau \kappa \delta$ such that the following holds. Assume that Z and W are independent of τ and that*

$$\lambda \geq (c_0/\sqrt{\delta}) \max\{\|Z\|_{L^2}, \|W\|_{L^2}\}.$$

If $(\tau_i)_{i=1}^m, (Z_i)_{i=1}^m$ and $(W_i)_{i=1}^m$ are independent copies of τ, Z and W respectively, then with probability at least

$$1 - 2 \exp(-c_1 m \delta) - 2 \exp(-c_2 m \|Z - W\|_{L^1}),$$

we have

$$|\{i : \text{sign}(Z_i + \tau_i) \neq \text{sign}(W_i + \tau_i)\}| \geq c_3 m \|Z - W\|_{L^1}.$$

Proof. Let θ be a constant to be specified later and observe that $\mathbb{P}(|Z| \geq \|Z\|_{L^2}/\sqrt{\theta}) \leq \theta$. Hence, with probability at least $1 - 2 \exp(-c_1 \theta m)$,

$$|\{i : |Z_i| \geq \|Z\|_{L^2}/\sqrt{\theta}\}| \leq 2\theta m,$$

where c_1 is an absolute constant; a similar estimate holds for $(W_i)_{i=1}^m$.

At the same time, recall that $\mathbb{P}(|Z - W| \geq \kappa \|Z - W\|_{L^1}) \geq \delta$, implying that with probability at least $1 - 2 \exp(-c_2 \delta m)$,

$$|\{i : |Z_i - W_i| \geq \kappa \|Z - W\|_{L^1}\}| \geq \delta m/2.$$

Set $\theta = \delta/16$ and let $\lambda \geq 4 \max\{\|Z\|_{L^2}/\sqrt{\delta}, \|W\|_{L^2}/\sqrt{\delta}\}$. The above shows that there is an event \mathcal{A} of (Z, W) -probability at least $1 - 2 \exp(-c_3 \delta m)$ on which the following holds: there exists $J \subset [m]$ of cardinality at least $\delta m/4$ such that for every $j \in J$,

$$|Z_j| \leq \lambda, \quad |W_j| \leq \lambda, \quad |Z_j - W_j| \geq \kappa \|Z - W\|_{L^1}.$$

Now fix two sequences of numbers $(z_i)_{i=1}^m$ and $(w_i)_{i=1}^m$ and consider the independent events

$$E_i = \{\text{sign}(z_i + \tau_i) \neq \text{sign}(w_i + \tau_i)\}, \quad 1 \leq i \leq m.$$

Recall that by (2.6), for every $i \in [m]$, if $|z_i| \leq \lambda$ and $|w_i| \leq \lambda$ then

$$\begin{aligned} \mathbb{P}_\tau(\text{sign}(z_i + \tau_i) \neq \text{sign}(w_i + \tau_i)) &= \mathbb{P}_\tau(z_i + \tau_i > 0, w_i + \tau_i \leq 0) + \mathbb{P}_\tau(z_i + \tau_i \leq 0, w_i + \tau_i > 0) \\ &= \mathbb{P}_\tau(-z_i < \tau \leq -w_i) + \mathbb{P}_\tau(-w_i < \tau \leq -z_i) \\ &\geq c_\tau |z_i - w_i|. \end{aligned}$$

Hence, for every realization of $(Z_i)_{i=1}^m$ and $(W_i)_{i=1}^m$ from the event \mathcal{A} ,

$$|\{j : \mathbb{P}_\tau(E_j) \geq c_\tau \kappa \|Z - W\|_{L^1}\}| \geq \delta m/4.$$

It follows that there are absolute constants c_4 and c_5 such that with τ -probability at least $1 - 2 \exp(-c_4 c_\tau \kappa \delta m \|Z - W\|_{L^1})$,

$$\sum_{i=1}^m \mathbb{1}_{E_i} \geq \sum_{j \in J} \mathbb{1}_{E_j} \geq \frac{|J|}{2} \cdot c_\tau \kappa \|Z - W\|_{L^1} \geq c_5 c_\tau \kappa \delta m \|Z - W\|_{L^1}.$$

Thus, with the wanted probability with respect to $(Z_i)_{i=1}^m, (W_i)_{i=1}^m$ and $(\tau_i)_{i=1}^m$, one has

$$|\{i : \text{sign}(Z_i + \tau_i) \neq \text{sign}(W_i + \tau_i)\}| \geq c_5 c_\tau \kappa \delta m \|Z - W\|_{L^1},$$

as claimed. □

Next, let us consider the random variable τ and the random vector X from Theorem 2.2: $\tau \sim \mathcal{U}[-\lambda, \lambda]$ and X is isotropic, symmetric and satisfies an L^1 - L^2 norm equivalence

with constant L . By the Paley–Zygmund inequality (see, e.g., [6]) there are constants κ and δ that depend only on L for which, for every $t \in \mathbb{R}^n$,

$$\mathbb{P}(|\langle X, t \rangle| \geq \kappa \|\langle X, t \rangle\|_{L^1}) \geq \delta.$$

Therefore, τ satisfies (2.6) with constant $c_\tau = 1/(2\lambda)$ and the random variables $Z = \langle X, x \rangle$ and $W = \langle X, w \rangle$ satisfy Lemma 2.6 with constants κ and δ that depend only on the equivalence constant L .

Proof of Theorem 2.5. Clearly, by Lemma 2.6,

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \geq c(L)m\|x - y\|_2/\lambda$$

with the promised probability, using the fact that

$$\max \{\|Z\|_{L^2}, \|W\|_{L^2}\} = \max \{\|\langle X, x \rangle\|_{L^2}, \|\langle X, y \rangle\|_{L^2}\} \leq R.$$

One has to show that in addition, $|\langle X_i, x \rangle + \tau_i|$ and $|\langle X_i, y \rangle + \tau_i|$ are also reasonably large. To that end, one may apply (2.4) twice, for $Z = \langle X, x \rangle$ and $Z = \langle X, y \rangle$, to see that for any $\varepsilon > 0$,

$$\max \{\mathbb{P}(|\langle X, x \rangle + \tau| \leq \varepsilon), \mathbb{P}(|\langle X, y \rangle + \tau| \leq \varepsilon)\} \leq \varepsilon/\lambda.$$

Therefore, with probability at least $1 - 2 \exp(-c\varepsilon m/\lambda)$, there are at most $4\varepsilon m/\lambda$ indices i for which

$$\min \{|\langle X_i, x \rangle + \tau|, |\langle X_i, y \rangle + \tau|\} \leq \varepsilon;$$

hence, setting $\varepsilon = (c(L)/8)\|x - y\|_2$ completes the proof. □

Next, one has to use the individual high probability estimate from Theorem 2.5 to obtain a uniform estimate in T . The idea is to use a covering argument combined with a simple stability property:

Lemma 2.7. *Fix a realization of X and τ and fix $r' > 0$. Assume that $\|w - v\|_2 \geq r'$ and that*

$$|\langle X, x - v \rangle| \leq \theta r'/3, \quad |\langle X, y - w \rangle| \leq \theta r'/3.$$

If v and w are θ -well-separated by $H_{X,\tau}$ then x and y are separated by $H_{X,\tau}$.

Proof. Since v and w are θ -well-separated by $H_{X,\tau}$, one has

$$\text{sign}(\langle X, v \rangle + \tau) \neq \text{sign}(\langle X, w \rangle + \tau), \quad |\langle X, v \rangle + \tau| \geq \theta \|v - w\|_2, \quad |\langle X, w \rangle + \tau| \geq \theta \|v - w\|_2.$$

Therefore, if

$$|\langle X, x - v \rangle| \leq \theta r'/3 \quad \text{and} \quad |\langle X, y - w \rangle| \leq \theta r'/3$$

it follows that $\text{sign}(\langle X, x \rangle + \tau) \neq \text{sign}(\langle X, y \rangle + \tau)$ (see Figure 6 for an illustration). □




Fig. 6. A ‘good’ hyperplane H_{X_i, τ_i} that well-separates v and w . On the one hand, one needs to shift the hyperplane in parallel by a distance proportional to $\theta \|v - w\|_2$ to hit w (the shift is marked in red). On the other hand, the parallel shift needed to hit y when starting from w is less than half this distance (shift marked in blue). As a consequence, a good hyperplane separates x and y .

The key component in the proof of Theorem 2.2 is the following fact:

Theorem 2.8. *There exist constants c_0, \dots, c_6 that depend only on L for which the following holds. Let $\lambda \geq c_0 R$, $r' \leq \lambda/2$, and $r'' \leq r'/4$. Assume that*

$$\log \mathcal{N}(T, r'') \leq c_1 m r' / \lambda, \tag{2.7}$$

and that

$$\mathbb{E} \sup_{z \in (T-T) \cap r'' B_2^n} |\{i : |\langle X_i, z \rangle| \geq c_2 r'\}| \leq m c_3 r' / \lambda. \tag{2.8}$$

Then with probability at least $1 - 8 \exp(-c_4 m (r'/\lambda)^2)$, for every $x, y \in T$ such that $\|x - y\|_2 \geq 2r'$,

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \geq c_5 m r' / \lambda,$$

and for every $x, y \in T$ such that $\|x - y\|_2 \leq r''/2$,

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \leq c_6 m r' / \lambda. \tag{2.9}$$

Proof. Let $V \subset T$ be an r' -cover of T . We apply (2.5) to every $Z = \langle X, v \rangle$, $v \in V$, and Theorem 2.5 to every pair of points from V . Let $c_1 \leq \min\{c, c_2\}/2$, where c and c_2 are as in the probability estimates of (2.5) and Theorem 2.5, respectively. If

$$\log |V| \leq c_1 m r' / \lambda$$

then by the union bound there is an event \mathcal{A}_1 of probability at least $1 - 6 \exp(-c_2 m r' / \lambda)$ such that for every $v \in V$,

$$|\{i : |\langle X_i, v \rangle + \tau_i| \geq r'\}| \geq \left(1 - \frac{2r'}{\lambda}\right) m \tag{2.10}$$

and if $v, w \in V$ satisfy $\|v - w\|_2 \geq r'$ then

$$|I_{v,w}(c_3)| \geq c_4 m \|v - w\|_2 / \lambda,$$

where the constants c_2, c_3 and c_4 depend only on L .

Now fix $x, y \in T$ that satisfy $\|x - y\|_2 \geq 2r'$ and let v, w be the nearest points in V to x and y respectively. By Lemma 2.7, if $i \in I_{v,w}(c_3)$ and

$$|\langle X_i, x - v \rangle|, |\langle X_i, y - w \rangle| \leq \frac{c_3}{3} r',$$

then x and y are separated by H_{X_i, τ_i} .

Note that $x - v, y - w \in (T - T) \cap r'' B_2^n$, let $\tilde{c}_3 = \min\{c_3, 1\}$ and set \mathcal{A}_2 to be the event

$$\sup_{z \in (T - T) \cap r'' B_2^n} |\{i : |\langle X_i, z \rangle| \geq (\tilde{c}_3/3)r'\}| \leq \frac{c_4}{2} \cdot \frac{mr'}{\lambda}. \tag{2.11}$$

Hence, on $\mathcal{A}_1 \cap \mathcal{A}_2$, if $\|x - y\|_2 \geq r'$ then

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \geq \frac{c_4}{2} \cdot \frac{mr'}{\lambda},$$

which is the wanted lower bound.

At the same time, if $\|x - v\|_2 \leq r''$ then by combining (2.10) and (2.11), one has the upper bound

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, v \rangle + \tau_i)\}| \leq c_5 \frac{mr'}{\lambda}.$$

All that is left is to estimate the probability of the event \mathcal{A}_2 . Note that

$$|\{i \in [m] : |\langle X_i, w \rangle| \geq (c_3/3)r'\}| = \sum_{i=1}^m \mathbb{1}_{\{|\langle X_i, w \rangle| \geq (c_3/3)r''\}} =: H_w,$$

and by the bounded differences inequality (see e.g. [4, Theorem 6.2]),

$$\mathbb{P}\left(\sup_{w \in (T - T) \cap r'' B_2^n} H_w \geq \mathbb{E} \sup_{w \in (T - T) \cap r'' B_2^n} H_w + mt\right) \leq 2 \exp(-cmt^2)$$

for a suitable absolute constant c . The claim follows with the choice of $t = (c_4/4) \cdot (r'/\lambda)$. □

Proof of Theorem 2.2. We apply Theorem 2.8 for the choice $r' = \rho/2$. Let us identify the conditions on r'' one has to impose to ensure that (2.8) is satisfied.

By the Giné–Zinn symmetrization theorem [11] and the contraction inequality for Bernoulli processes [16], one has

$$\begin{aligned} \mathbb{E} \sup_{z \in (T - T) \cap r'' B_2^n} |\{i : |\langle X_i, z \rangle| \geq c\rho\}| &\leq \mathbb{E} \sup_{z \in (T - T) \cap r'' B_2^n} \frac{1}{c\rho} \sum_{i=1}^m |\langle X_i, z \rangle| \\ &\leq \mathbb{E} \sup_{z \in (T - T) \cap r'' B_2^n} \frac{2}{c\rho} \left| \sum_{i=1}^m \varepsilon_i \langle X_i, z \rangle \right| + \frac{mr''}{c\rho} =: (1) + (2). \end{aligned}$$

To satisfy (2.8) it suffices to bound both terms by $cm\rho/\lambda$. The required estimate on (2) holds once

$$r'' \leq c(L)\rho^2/\lambda,$$

and to ensure a suitable estimate on (1) it suffices that

$$E(T_{r''}) \leq c(L)\sqrt{m}\rho^2/\lambda.$$

The claim follows by setting $r = r''$.

This immediately yields the lower bound in Theorem 2.2. To complete the proof of the upper bound, recall that T is (r, γ) -metrically convex. For given x, y with $\|x - y\|_2 \geq 2r$ let $(z_j)_{j=1}^\ell$ be as in Definition 2.1. Then

$$\begin{aligned} & |\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \\ & \leq \sum_{j=0}^\ell |\{i : \text{sign}(\langle X_i, z_j \rangle + \tau_i) \neq \text{sign}(\langle X_i, z_{j+1} \rangle + \tau_i)\}|, \end{aligned}$$

and the claim follows from the ‘local’ upper bound (2.9). □

2.2. The subgaussian scenario

When X is an L -subgaussian random vector one may establish an improved version of Theorem 2.8: first, by showing that one may take r'' to be of the order of r' up to a logarithmic factor; and second, by providing a better probability estimate on the outcome. Moreover, thanks to the subgaussian property, one may replace the empirical parameter $E(T_r)$ by its Gaussian counterpart, $\ell_*(T_r)$.

Theorem 2.9. *There exist constants c_0, \dots, c_5 that depend only on L for which the following holds. Assume that $\lambda \geq c_0R$; that*

$$r'' \leq c_1 \frac{r'}{\sqrt{\log(e\lambda/r')}};$$

that

$$\log \mathcal{N}(T, r'') \leq c_2 \frac{mr'}{\lambda}; \tag{2.12}$$

and that

$$\ell_*(T_{r''}) \leq c_3 \sqrt{m} \frac{(r')^{3/2}}{\sqrt{\lambda}}. \tag{2.13}$$

Then with probability at least $1 - 8 \exp(-c_4mr'/\lambda)$, for every $x, y \in T$ such that $\|x - y\|_2 \geq 2r'$,

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \geq c_5 \frac{mr'}{\lambda},$$

and for every $x, y \in T$ such that $\|x - y\|_2 \leq r''/2$,

$$|\{i : \text{sign}(\langle X_i, x \rangle + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \leq c_6 \frac{mr'}{\lambda}.$$

The only difference between the proof of Theorem 2.9 and that of Theorem 2.8 is the control one has on the probability that

$$\sup_{z \in (T-T) \cap r''B_2^d} |\{i : |\langle X_i, z \rangle| \geq cr'\}| \leq Cmr'/\lambda. \tag{2.14}$$

When X merely satisfies an L^1 - L^2 norm equivalence, one has to resort to the bounded differences inequality for a high probability estimate. However, when X is L -subgaussian there is more information on the behaviour of the supremum. Specifically, we use the following fact.

Theorem 2.10. *Let X be an isotropic L -subgaussian random vector and let $S \subset \mathbb{R}^n$. If $1 \leq k \leq m$ and $u \geq 1$ then with probability at least $1 - 2 \exp(-c_1 u^2 k \log(em/k))$,*

$$\sup_{z \in S} \max_{|I| \leq k} \left(\sum_{i \in I} \langle X_i, z \rangle^2 \right)^{1/2} \leq c_2 (\ell_*(S) + u d_S \sqrt{k \log(em/k)}),$$

where c_1 and c_2 depend only on L and $d_S = \sup_{z \in S} \|z\|_2$.

We omit the proof of Theorem 2.10, which is standard. It is based on generic chaining (see e.g. [7, Theorem 3.2]) combined with Talagrand’s majorizing measures theorem [25].

Proof of Theorem 2.9. Observe that if $(a_i)_{i=1}^m$ is a sequence of nonnegative numbers then the k -largest element satisfies

$$a_k^* \leq \max_{|I| \leq k} \left(\frac{1}{k} \sum_{i \in I} a_i^2 \right)^{1/2}.$$

With that in mind, one has to ensure that for $k = Cmr'/\lambda$,

$$\sup_{z \in (T-T) \cap r'' B_2^n} \max_{|I| \leq k} \left(\frac{1}{k} \sum_{i \in I} \langle X_i, z \rangle^2 \right)^{1/2} \leq cr',$$

and by Theorem 2.10, it suffices to verify that

$$\left(\frac{\sqrt{\lambda} \ell_*(T_{r''})}{\sqrt{mr'}} + r'' \sqrt{\log(e\lambda/r')} \right) \leq c_1(L)r'.$$

Clearly, the wanted estimate follows if

$$r'' \leq c_2(L) \frac{r'}{\sqrt{\log(e\lambda/r')}} \quad \text{and} \quad m \geq c_3(L) \lambda \frac{(\ell_*(T_{r''}))^2}{(r')^3};$$

and in that case, (2.14) holds with probability at least $1 - 2 \exp(-c_4 m (r'/\lambda) \log(e\lambda/r'))$.

The rest of the proof of Theorem 2.9 is identical to that of Theorem 2.8 and is omitted. □

Now one may complete the proof of Theorem 2.3, by setting $\rho = 2r'$ and $r = r''/2$, and noting that Theorem 2.9 yields the lower bound and the ‘local’ upper bound. The upper bound follows directly from the local upper bound and the metric convexity assumption (see the end of the proof of Theorem 2.2 for this argument). □

3. Random tessellations: noisy measurements

With the machinery developed for the proofs of Theorems 2.2 and 2.3 at our disposal, let us present the proofs of Theorems 1.9 and 1.11.

Recall that $T \subset RB_2^n$ and that X is an isotropic and symmetric random vector, while $\tau \sim \mathcal{U}[-\lambda, \lambda]$. The rows of the measurement matrix A are $(X_i)_{i=1}^m$ and the given observations are the coordinates of the vector q_{corr} , which is a corrupted version of

$$\text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}) = (\langle X_i, x \rangle + v_i + \tau_i)_{i=1}^m,$$

by at most βm ‘sign flips’. In the first scenario, X and v satisfy an L^1 - L^2 norm equivalence with constant L , while in the second they are L -subgaussian.

The goal is to show that there is a constant C depending only on λ and L so that with high probability, for any $x, y \in T$ that satisfy $\|x - y\|_2 \geq \rho$,

$$|\{i : \text{sign}(\langle X_i, x \rangle + v_i + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)\}| \geq Cm\|x - y\|_2, \tag{3.1}$$

and at the same time,

$$|\{i : \text{sign}(\langle X_i, x \rangle + v_i + \tau_i) \neq \text{sign}(\langle X_i, x \rangle + \tau_i)\}| < Cm\rho/4. \tag{3.2}$$

Together these conditions imply that the recovery program (1.9), which minimizes the Hamming distance between q_{corr} and $(\text{sign}(\langle X_i, z \rangle + \tau_i))_{i=1}^m$ with respect to $z \in T$, achieves reconstruction accuracy ρ as long as the fraction of the corrupted bits is at most $\beta \leq C\rho/4$. Indeed, (3.2) implies that any solution $x^\#$ of (1.9) must satisfy

$$d_H(q_{\text{corr}}, (\text{sign}(\langle X_i, x^\# \rangle + \tau_i))_{i=1}^m) < Cm\rho/2$$

and then (3.1) shows that $\|x^\# - x\|_2 \leq \rho$.

The proofs of (3.1) in both scenarios follow from minor modifications of the results established in the previous section. Rather than repeating the arguments, let us sketch the adjustments one has to make.

First, one has to consider a modified notion of being ‘well-separated’ by a hyperplane:

Definition 3.1. The hyperplane H_{X_i, τ_i} θ -well-separates x and y if

- $\text{sign}(\langle X_i, x \rangle + v_i + \tau_i) \neq \text{sign}(\langle X_i, y \rangle + \tau_i)$,
- $|\langle X_i, x \rangle + v_i + \tau_i| \geq \theta\|x - y\|_2$, and
- $|\langle X_i, y \rangle + \tau_i| \geq \theta\|x - y\|_2$.

Denote by $J_{x,y}(\theta) \subset [m]$ the set of indices for which H_{X_i, τ_i} θ -well-separates x and y .

Next, one has to establish the analog of Theorem 2.5 and show that every pair x, y is well-separated by a fraction of the hyperplanes that is proportional to $\|x - y\|_2$.

Theorem 3.2. *There are constants c_1, \dots, c_4 that depend only on L for which the following holds. Let $x, y \in RB_2^n$ and fix $\lambda \geq c_1 \max\{R, \|v\|_{L^2}\}$. With probability at least $1 - 4 \exp(-c_2 m\|x - y\|_2/\lambda)$,*

$$|J_{x,y}(c_3)| \geq c_4 m\|x - y\|_2/\lambda.$$

Just as in Theorem 2.5, the proof is an outcome of Lemma 2.6, only this time with the choice $Z = \langle X, x \rangle + v$ and $W = \langle X, y \rangle$. To that end, one has to verify that $Z - W$ satisfies a small-ball condition, and that is immediate from the following observation and the Paley–Zygmund inequality.

Lemma 3.3. *Let Z and W be as above. Then*

$$\|Z - W\|_{L^1} \geq \frac{1}{2L} \|Z - W\|_{L^2}.$$

Proof. Let $\alpha, \beta \in \mathbb{R}$ and set ε to be a symmetric, $\{-1, 1\}$ -valued random variable. Clearly,

$$\mathbb{E}_\varepsilon |\varepsilon\alpha + \beta| \geq \frac{1}{2}(|\alpha| + |\beta|).$$

Since X is a symmetric random vector, $\langle X, x - y \rangle$ has the same distribution as $\varepsilon\langle X, x - y \rangle$, where ε is independent of X and of v . Hence,

$$\begin{aligned} \|Z - W\|_{L^1} &= \|\langle X, x - y \rangle + v\|_{L^1} = \mathbb{E}_{X,v} \mathbb{E}_\varepsilon |\varepsilon\langle X, x - y \rangle + v| \\ &\geq \frac{1}{2}(\mathbb{E}|\langle X, x - y \rangle| + \mathbb{E}|v|) \geq \frac{1}{2L}(\| \langle X, x - y \rangle \|_{L^2} + \|v\|_{L^2}) \\ &\geq \frac{1}{2L} \| \langle X, x - y \rangle + v \|_{L^2}. \quad \square \end{aligned}$$

The other components needed to ensure separation in the sense of Definition 3.1 follow from an identical argument used in the proof of Theorem 2.5, by conditioning on X_i and v_i rather than just on X_i .

Theorem 3.2 allows one to control the set V of ‘centres’ of a cover of T , and all that remains now is to show that if x' is close to a centre x and y' is close y then there will be few indices for which

$$\text{sign}(\langle X_i, x \rangle + v_i + \tau_i) \neq \text{sign}(\langle X_i, x' \rangle + v_i + \tau_i)$$

or

$$\text{sign}(\langle X_i, y \rangle + \tau_i) \neq \text{sign}(\langle X_i, y' \rangle + \tau_i).$$

In both cases, and using the notation of the previous section, one may follow the argument in the proof of Theorem 2.8. Thus, it suffices to show that

$$\sup_{z \in (T-T) \cap r'' B_2^d} |\{i : |\langle X_i, z \rangle| \geq c_1 r'\}| \leq c_2 m r' / \lambda, \tag{3.3}$$

and that concludes the proof of the bound (3.1) with $C \sim_L 1/\lambda$. □

Let us turn to (3.2). In the heavy-tailed case, one may invoke the proof of Theorem 2.8 to show that with probability at least $1 - 8 \exp(-cm(\rho/\lambda)^2)$, for every $x \in T$,

$$|\{i : |\langle X_i, x \rangle + \tau_i| \geq \rho/2\}| \geq m \left(1 - \frac{\rho}{\lambda}\right).$$

If $|\mathbb{E}v| \leq \rho/16$ and $\sigma^2 \leq (1/64)\rho^3/\lambda$, it follows that

$$\mathbb{P}(|v| \geq \rho/4) \leq \mathbb{P}(|v - \mathbb{E}v| \geq \rho/8) \leq \rho/\lambda.$$

Hence, with probability at least $1 - 2 \exp(-cm\rho/\lambda)$,

$$|\{i : |v_i| \geq \rho/4\}| \leq 2\rho m/\lambda.$$

On the intersection of the two events, for every $x \in T$ one has

$$|\{i : \text{sign}(\langle X_i, x \rangle + v_i + \tau_i) \neq \text{sign}(\langle X_i, x \rangle + \tau_i)\}| \leq 3\rho m/\lambda,$$

as required.

The proof in the subgaussian case is analogous and therefore omitted. □

4. Robust recovery via a convex program

This section is devoted to the proof of Theorem 1.13. Set

$$\phi(z) = \frac{1}{m} \langle q_{\text{corr}}, Az \rangle - \frac{1}{2\lambda} \|z\|_2^2, \tag{4.1}$$

then the convex optimization procedure (1.12) is exactly

$$\max_{z \in \text{conv}(T)} \phi(z).$$

Recall that $U = \text{conv}(T)$ and that $U_\rho = (U - U) \cap \rho B_2^n$; X_1, \dots, X_m are the rows of the matrix A which are distributed according to an isotropic, symmetric, L -subgaussian random vector X ; and $\tau \sim \mathcal{U}[-\lambda, \lambda]$. Here we assume for the sake of simplicity that v has mean zero and variance σ^2 , though the modifications needed to handle the case in which v has a nontrivial adversarial component are straightforward. Finally, as before $q_{\text{corr}} \in \{-1, 1\}^m$ satisfies

$$d_H(q_{\text{corr}}, \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}})) \leq \beta m. \tag{4.2}$$

As in most regularized procedures, the idea is to study the ‘excess functional’ $\phi(z) - \phi(x)$: for a reconstruction error ρ , we determine a sufficient condition on the number of measurements m which guarantees that $\phi(z) - \phi(x) < 0$ whenever $z \in U$ and $\|x - z\|_2 \geq \rho$. Clearly, that implies that the solution $x^\#$ to (4.1) satisfies $\|x^\# - x\|_2 \leq \rho$. As before, the goal is to obtain a uniform estimate, i.e., the high probability event for which the above holds should not depend on the identity of $x \in T$.

The first step towards a uniform estimate is a decomposition of the excess functional. Note that

$$\begin{aligned} \phi(z) - \phi(x) &= \frac{1}{m} (\langle q_{\text{corr}}, Az \rangle - \langle q_{\text{corr}}, Ax \rangle) - \frac{1}{2\lambda} \|z\|_2^2 + \frac{1}{2\lambda} \|x\|_2^2 \\ &= \frac{1}{m} \langle q_{\text{corr}} - \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}), A(z - x) \rangle \\ &\quad + \frac{1}{m} (\langle \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}), A(z - x) \rangle - \mathbb{E} \langle \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}), A(z - x) \rangle) \\ &\quad + \frac{1}{m} \mathbb{E} \langle \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}), A(z - x) \rangle - \frac{1}{2\lambda} \|z\|_2^2 + \frac{1}{2\lambda} \|x\|_2^2 \\ &=: (1) + (2) + (3), \end{aligned}$$

We use this decomposition to find constants C and $\rho > 0$ and a high probability event on which, for every $x \in T$ and $z \in U$,

$$|(1)| \leq C\|x - z\|_2^2, \quad |(2)| \leq C\|x - z\|_2^2, \quad (3) \leq -4C\|x - z\|_2^2, \quad (4.3)$$

provided that $\|x - z\|_2 \geq \rho$.

Estimating (3). The starting point is a straightforward observation: for $\tau \sim \mathcal{U}[-\lambda, \lambda]$ and any $z \in \mathbb{R}$,

$$\mathbb{E} \operatorname{sign}(z + \tau) = \frac{z}{\lambda} \mathbb{1}_{\{|z| \leq \lambda\}} + \mathbb{1}_{\{z > \lambda\}} - \mathbb{1}_{\{z < -\lambda\}}. \quad (4.4)$$

Lemma 4.1. *There exist absolute constants C and c for which the following holds. Let Z and W be random variables and let $\tau \sim \mathcal{U}[-\lambda, \lambda]$ be independent of Z and W . Then*

$$\left| \mathbb{E} W \operatorname{sign}(Z + \tau) - \frac{1}{\lambda} \mathbb{E} W Z \right| \leq C \|W\|_{L_2} \max \left\{ 1, \frac{\|Z\|_{\psi_2}}{\lambda} \right\} \exp(-c\lambda^2 / \|Z\|_{\psi_2}^2).$$

Proof. By (4.4),

$$\begin{aligned} \mathbb{E}_\tau W \operatorname{sign}(Z + \tau) &= W \left(\frac{Z}{\lambda} \mathbb{1}_{\{|Z| \leq \lambda\}} + \mathbb{1}_{\{Z > \lambda\}} - \mathbb{1}_{\{Z < -\lambda\}} \right) \\ &= W \left(\frac{Z}{\lambda} - \frac{Z}{\lambda} \mathbb{1}_{\{|Z| > \lambda\}} + \mathbb{1}_{\{Z > \lambda\}} - \mathbb{1}_{\{Z < -\lambda\}} \right) \\ &=: \frac{WZ}{\lambda} + (*). \end{aligned}$$

Hence, $\mathbb{E} \operatorname{sign}(Z + \tau) W = \frac{1}{\lambda} \mathbb{E} W Z + \mathbb{E}(*),$ and all that is left to show is

$$\mathbb{E}|(*)| \leq C \|W\|_{L_2} \max \left\{ 1, \frac{\|Z\|_{\psi_2}}{\lambda} \right\} \exp(-c\lambda^2 / \|Z\|_{\psi_2}^2)$$

for absolute constants C and c .

Note that $\mathbb{E}|WZ \mathbb{1}_{\{|Z| > \lambda\}}| \leq \|W\|_{L_2} (\mathbb{E} Z^2 \mathbb{1}_{\{|Z| > \lambda\}})^{1/2}$. By tail integration,

$$\begin{aligned} \mathbb{E} Z^2 \mathbb{1}_{\{|Z| > \lambda\}} &\leq \lambda^2 \mathbb{P}(|Z| > \lambda) + 2 \int_\lambda^\infty t \mathbb{P}(|Z| > t) dt \\ &\leq (\lambda^2 + \|Z\|_{\psi_2}^2) \exp(-c_1 \lambda^2 / \|Z\|_{\psi_2}^2), \end{aligned}$$

where c_1 is a suitable absolute constant. The estimate on the other two terms follows because $\mathbb{1}_{\{|Z| > \lambda\}} \leq (|Z|/\lambda) \mathbb{1}_{\{|Z| > \lambda\}}$. □

Corollary 4.2. *There exist absolute constants c and C for which the following holds. For every $x, z \in \mathbb{R}^n$,*

$$\begin{aligned} &\frac{1}{m} \mathbb{E} \langle \operatorname{sign}(Ax + \nu_{\text{noise}} + \tau_{\text{thres}}), A(z - x) \rangle \\ &\leq \frac{1}{\lambda} \langle x, z - x \rangle + C \|z - x\|_2 \max \left\{ 1, \frac{L(\sigma + \|x\|_2)}{\lambda} \right\} \exp \left(-c \frac{\lambda^2}{L^2(\|x\|_2^2 + \sigma^2)} \right), \end{aligned}$$

where, as always, L is the subgaussian constant of X and of v . In particular,

$$\begin{aligned} & \frac{1}{m} \mathbb{E} \langle \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}), A(z - x) \rangle - \frac{1}{2\lambda} \|z\|_2^2 + \frac{1}{2\lambda} \|x\|_2^2 \\ & \leq -\frac{1}{2\lambda} \|z - x\|_2^2 + C \|z - x\|_2 \max \left\{ 1, \frac{L(\sigma + R)}{\lambda} \right\} \exp \left(-c \frac{\lambda^2}{L^2(R^2 + \sigma^2)} \right). \end{aligned}$$

Proof. The first statement follows from Lemma 4.1 for the choices of $Z_i = \langle X_i, x \rangle + v_i$ and $W_i = \langle X_i, z - x \rangle$: recalling that the v_i are centred, have variance σ^2 and are independent of X_i ,

$$\|Z_i\|_{\psi_2}^2 \leq c(\| \langle X_i, x \rangle \|_{\psi_2}^2 + \|v_i\|_{\psi_2}^2) \leq cL^2(\|x\|_2^2 + \sigma^2) \leq cL^2(R^2 + \sigma^2);$$

and, because X is isotropic,

$$\| \langle X_i, z - x \rangle \|_{L_2} = \|z - x\|_2 \quad \text{and} \quad \frac{1}{\lambda} \mathbb{E}(\langle X_i, x \rangle + v_i) \langle X_i, z - x \rangle = \frac{1}{\lambda} \langle x, z - x \rangle.$$

The ‘in particular’ part is evident because

$$\frac{1}{\lambda} \langle x, z - x \rangle - \frac{1}{2\lambda} \|z\|_2^2 + \frac{1}{2\lambda} \|x\|_2^2 = -\frac{1}{2\lambda} \|z - x\|_2^2. \quad \square$$

Corollary 4.2 leads to the wanted estimate on (3):

Corollary 4.3. *There are constants c_1 and c_2 that depend only on L such that if*

$$\lambda \geq c_1(\sigma + R)\sqrt{\log(c_2/\rho)} \tag{4.5}$$

and $\|x - z\|_2 \geq \rho$ then

$$\frac{1}{m} \mathbb{E} \langle \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}), A(z - x) \rangle - \frac{1}{2\lambda} \|z\|_2^2 + \frac{1}{2\lambda} \|x\|_2^2 \leq -\frac{1}{4\lambda} \|z - x\|_2^2. \tag{4.6}$$

Estimating (1). Next, let us estimate |(1)| from above, by studying

$$\sup_{x \in T} \sup_{\{z \in U: \|z-x\|_2 \geq \rho\}} \left| \frac{1}{m} \langle q_{\text{corr}} - \text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}), A(z - x) / \|z - x\|_2 \rangle \right| =: (*).$$

Observe that for every x, z , and $\eta_i = (q_{\text{corr}})_i - \text{sign}(\langle X_i, x \rangle + v_i + \tau_i)$, one has

$$\left| \frac{1}{m} \sum_{i=1}^m \eta_i \langle X_i, z - x \rangle \right| \leq \frac{2}{m} \sum_{\{i: \eta_i \neq 0\}} |\langle X_i, z - x \rangle|$$

and $(\eta_i)_{i=1}^m$ has at most βm nonzero coordinates. Hence, for any X_1, \dots, X_m ,

$$(*) \leq \max_{|I| \leq \beta m} \sup_{\{x, z \in U: \|z-x\|_2 \geq \rho\}} \frac{1}{m} \sum_{i \in I} |\langle X_i, (z - x) / \|z - x\|_2 \rangle|.$$

Therefore, taking into account the estimate (4.6) on (3), it suffices to show that for every $x, z \in U$ such that $\|x - z\|_2 \geq \rho$,

$$\max_{|I| \leq \beta m} \frac{1}{m} \sum_{i \in I} |\langle X_i, (z - x) / \|z - x\|_2^2 \rangle| \leq \frac{1}{16\lambda}.$$

To that end, observe that if $f : \mathbb{R}^n \rightarrow \mathbb{R}_+$ is positive homogeneous and $W \subset \mathbb{R}^n$ is star-shaped around 0, i.e., $\theta w \in W$ for all $w \in W$ and $0 < \theta < 1$, then

$$\sup_{\{w \in W : \|w\|_2 \geq \rho\}} f(w / \|w\|_2^2) = \sup_{\{w \in W : \|w\|_2 = \rho\}} f(w) / \rho^2.$$

We will refer to this argument, which reflects the general fact that star-shaped sets become ‘relatively richer’ close to their centres, as a ‘star-shape argument’.

Theorem 4.4. *There exist constants c_1, c_2 , and c_3 depending only on L for which the following holds. Assume that*

$$\rho \geq c_1 \lambda \beta \sqrt{\log(e/\beta)} \quad \text{and} \quad \ell_*(U_\rho) \leq c_2 \sqrt{\frac{m}{\beta}} \cdot \frac{\rho^2}{\lambda}. \tag{4.7}$$

Then with probability at least $1 - 2 \exp(-c_3 \beta m \log(e/\beta))$, for every $x, z \in U$ such that $\|z - x\|_2 \geq \rho$ one has

$$\max_{|I| \leq \beta m} \frac{1}{m} \sum_{i \in I} |\langle X_i, z - x \rangle| \leq \frac{1}{16\lambda} \|z - x\|_2^2.$$

Proof. Since $U - U$ is star-shaped around 0, the star-shape argument implies that for any $I \subset [m]$,

$$\sup_{\{x, z \in U : \|z - x\|_2 \geq \rho\}} \frac{1}{m} \sum_{i \in I} \left| \left\langle X_i, \frac{z - x}{\|z - x\|_2^2} \right\rangle \right| = \sup_{w \in (U - U) \cap \rho S^{n-1}} \frac{1}{m \rho^2} \sum_{i \in I} |\langle X_i, w \rangle|.$$

Apply Theorem 2.10 for $k = \beta m$ and note that $\|w\|_1 \leq \sqrt{k} \|w\|_2$ for $w \in \mathbb{R}^k$. This shows that, with probability at least $1 - 2 \exp(-c(L) \beta m \log(e/\beta))$,

$$\sup_{w \in (U - U) \cap \rho S^{n-1}} \frac{1}{m \rho^2} \sum_{i \in I} |\langle X_i, w \rangle| \leq \frac{C(L)}{m \rho^2} (\sqrt{\beta m} \ell_*(U_\rho) + \rho \beta m \sqrt{\log(e/\beta)}) \leq \frac{1}{16\lambda},$$

provided that (4.7) holds. □

Estimating (2). Finally, let us derive an upper estimate on

$$\left| \frac{1}{m} \left((\text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}), A(z - x)) - \mathbb{E}(\text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}), A(z - x)) \right) \right|$$

that holds uniformly for all $x \in T$ and $z \in U$ such that $\|x - z\|_2 \geq \rho$.

Theorem 4.5. *There exist constants c_1, \dots, c_4 that depend only on L for which the following holds. Let $\lambda \geq c_1$ and assume that*

$$\log \mathcal{N}(T, \rho \log^{-1}(e\lambda/\rho)) \leq c_2 m \cdot \frac{\rho^2}{\lambda^2} \quad \text{and} \quad \ell_*(U_\rho) \leq c_3 \frac{\rho^2}{\lambda} \sqrt{m}.$$

Then with probability at least $1 - 5 \exp(-c_4 m \rho^2 / \lambda^2)$, for every $x \in T$ and $z \in U$ such that $\|x - z\|_2 \geq \rho$,

$$|(2)| \leq \frac{\rho^2}{16\lambda}.$$

The proof of Theorem 4.5 is based on a covering argument. Let $V \subset T$ be an r -cover for a well-chosen r . A crucial part of the proof is to show that for every $v \in V$ the (random) set of sign patterns

$$\mathbb{S}_v := \{\text{sign}(Ax + v_{\text{noise}} + \tau_{\text{thres}}) : x \in U, \|x - v\|_2 \leq r\}$$

is relatively simple: it consists of small perturbations of $\text{sign}(Av + v_{\text{noise}} + \tau_{\text{thres}})$.

Since this observation is essentially the second part of Theorem 2.9, we formulate it in the way that it will be applied and omit its proof—which is almost identical to the proof of Theorem 2.9.

Lemma 4.6. *There are constants c_0, c_1 and c_2 that depend only on L for which the following holds. Let $0 < r' \leq \lambda/2$ and $r'' \leq c_0 r' / \sqrt{\log(e\lambda/r')}$. Assume that*

$$\log \mathcal{N}(T, r'') \leq c_1 \frac{r'}{\lambda} m, \tag{4.8}$$

and that

$$\ell_*(U_{r''}) \leq c_1 \frac{(r')^{3/2}}{\lambda^{1/2}} \sqrt{m}. \tag{4.9}$$

If V is a minimal r'' -cover of T , then with probability at least $1 - 2 \exp(-c_2 m r' / \lambda)$, for every $v \in V$,

$$\mathbb{S}_v \subset \text{sign}(Av + v_{\text{noise}} + \tau_{\text{thres}}) + 2\mathcal{Z},$$

where \mathcal{Z} is the set of all $\{-1, 0, 1\}$ -valued vectors in \mathbb{R}^m that have at most $3(r'/\lambda)m$ nonzero coordinates.

Proof of Theorem 4.5. By a standard symmetrization argument it suffices to estimate the probability that

$$\sup_{x \in T} \sup_{\{z \in U : \|z - x\|_2 \geq \rho\}} \left| \frac{1}{m} \sum_{i=1}^m \varepsilon_i \text{sign}(\langle X_i, x \rangle + v_i + \tau_i) \left\langle X_i, \frac{x - z}{\|x - z\|_2} \right\rangle \right| \leq \frac{1}{32\lambda},$$

where $(\varepsilon_i)_{i=1}^m$ are independent, symmetric, $\{-1, 1\}$ -valued random variables that are also independent of $(X_i, v_i, \tau_i)_{i=1}^m$.

Let r', r'' , and V be as in Lemma 4.6 and observe that

$$\begin{aligned} & \sup_{x \in T} \sup_{\{z \in U: \|z-x\|_2 \geq \rho\}} \left| \frac{1}{m} \sum_{i=1}^m \varepsilon_i \operatorname{sign}(\langle X_i, x \rangle + v_i + \tau_i) \left\langle X_i, \frac{x-z}{\|x-z\|_2} \right\rangle \right| \\ &= \max_{v \in V} \sup_{\{x \in T: \|x-v\|_2 \leq r''\}} \sup_{\{z \in U: \|z-x\|_2 \geq \rho\}} \left| \frac{1}{m} \sum_{i=1}^m \varepsilon_i \operatorname{sign}(\langle X_i, x \rangle + v_i + \tau_i) \left\langle X_i, \frac{x-z}{\|x-z\|_2} \right\rangle \right|. \end{aligned}$$

Let \mathcal{A} be the event from Lemma 4.6. Recall that $\mathbb{P}(\mathcal{A}) \geq 1 - 2 \exp(-cmr'/\lambda)$ and that on \mathcal{A} , for any $v \in V$ and $x \in T$ such that $\|x - v\|_2 \leq r''$, $\operatorname{sign}(\langle X_i, x \rangle + v_i + \tau_i)$ differs from $\operatorname{sign}(\langle X_i, v \rangle + v_i + \tau_i)$ on at most $(3r'/\lambda)m$ indices. Therefore, for every $v \in V$,

$$\begin{aligned} & \sup_{\{x \in T: \|x-v\|_2 \leq r''\}} \sup_{\{z \in U: \|z-x\|_2 \geq \rho\}} \left| \frac{1}{m} \sum_{i=1}^m \varepsilon_i \operatorname{sign}(\langle X_i, x \rangle + v_i + \tau_i) \left\langle X_i, \frac{x-z}{\|x-z\|_2} \right\rangle \right| \\ & \leq \sup_{\{x \in T: \|x-v\|_2 \leq r''\}} \sup_{\{z \in U: \|z-x\|_2 \geq \rho\}} \left| \frac{1}{m} \sum_{i=1}^m \varepsilon_i \operatorname{sign}(\langle X_i, v \rangle + v_i + \tau_i) \left\langle X_i, \frac{x-z}{\|x-z\|_2} \right\rangle \right| \\ & \quad + 2 \sup_{\{x \in T: \|x-v\|_2 \leq r''\}} \sup_{\{z \in U: \|z-x\|_2 \geq \rho\}} \max_{|I| \leq (3r'/\lambda)m} \frac{1}{m} \sum_{i \in I} \left| \left\langle X_i, \frac{x-z}{\|x-z\|_2} \right\rangle \right| \\ & =: (a)_v + (b)_v. \end{aligned}$$

Observe that both $(a)_v$ and $(b)_v$ are homogeneous in $x - z$; therefore, by a star-shape argument

$$(a)_v \leq \frac{1}{\rho^2} \sup_{\{z, x \in U: \|z-x\|_2 = \rho\}} \left| \frac{1}{m} \sum_{i=1}^m \varepsilon_i \operatorname{sign}(\langle X_i, v \rangle + v_i + \tau_i) \langle X_i, x-z \rangle \right|, \tag{4.10}$$

$$(b)_v \leq \frac{1}{\rho^2} \sup_{\{z, x \in U: \|z-x\|_2 = \rho\}} \max_{|I| \leq (3r'/\lambda)m} \frac{1}{m} \sum_{i \in I} |\langle X_i, x-z \rangle|. \tag{4.11}$$

Let us begin by estimating the right hand side of (4.10). For every fixed $v \in V$ and conditioned on $(X_i, v_i, \tau_i)_{i=1}^m$, this is the supremum of a Bernoulli process indexed by a set of the form

$$\left\{ \sum_{i=1}^m a_i w_i : w \in W \right\},$$

where $(a_i)_{i=1}^m$ is a fixed vector of signs and

$$W \subset \{(\langle X_i, u \rangle)_{i=1}^m : u \in U_\rho\}.$$

By the contraction inequality for Bernoulli processes [16] applied conditionally on $(X_i, v_i, \tau_i)_{i=1}^m$, it follows that

$$\mathbb{P}\left(\sup_{u \in U_\rho} \left| \sum_{i=1}^m \varepsilon_i a_i \langle X_i, u \rangle \right| \geq t\right) \leq 2\mathbb{P}\left(\sup_{u \in U_\rho} \left| \sum_{i=1}^m \varepsilon_i \langle X_i, u \rangle \right| \geq t\right).$$

Since $\frac{1}{\sqrt{m}} \sum_{i=1}^m \varepsilon_i X_i$ is cL -subgaussian, it follows by generic chaining (see e.g. [7, Theorem 3.2]) combined with Talagrand’s majorizing measures theorem [25] that

$$\left\| \sup_{u \in U_\rho} \frac{1}{\sqrt{m}} \sum_{i=1}^m \varepsilon_i \langle X_i, u \rangle \right\|_{L_p} \leq c(L) \left(\ell_*(U_\rho) + \sqrt{p} \sup_{u \in U_\rho} \|u\|_2 \right).$$

Hence, for any $t \geq 1$, with probability at least $1 - e^{-t}$,

$$\frac{1}{\rho^2} \sup_{u \in U_\rho} \left| \frac{1}{m} \sum_{i=1}^m \varepsilon_i \langle X_i, u \rangle \right| \leq \frac{c_0}{\rho^2} \left(\frac{\ell_*(U_\rho)}{\sqrt{m}} + \rho \sqrt{\frac{t}{m}} \right).$$

Taking $t = c_1 m \rho^2 / \lambda^2$, it follows that if

$$\ell_*(U_\rho) \lesssim_L \frac{\rho^2}{\lambda} \sqrt{m}, \tag{4.12}$$

then for any fixed $v \in V$,

$$(a)_v \leq \frac{1}{64\lambda} \quad \text{with probability at least } 1 - \exp(-c_1 m \rho^2 / \lambda^2). \tag{4.13}$$

By the union bound, (4.13) holds uniformly for all $v \in V$ as long as

$$\log |V| = \log \mathcal{N}(T, r'') \leq \frac{c_1}{2} \cdot \frac{m \rho^2}{\lambda^2}. \tag{4.14}$$

Next, observe that by (4.11),

$$\sup_{v \in V} (b)_v \leq \frac{1}{\rho^2} \sup_{u \in U_\rho} \max_{|I| \leq (3r'/\lambda)m} \frac{1}{m} \sum_{i \in I} |\langle X_i, u \rangle|.$$

By Theorem 2.10 for $k = 3r'm/\lambda$, with probability at least $1 - 2 \exp(-c_2(r'/\lambda) \cdot \log(e\lambda/r')m)$,

$$\frac{1}{\rho^2} \sup_{u \in U_\rho} \max_{|I| \leq (3r'/\lambda)m} \frac{1}{m} \sum_{i \in I} |\langle X_i, v \rangle| \lesssim_L \left(\sqrt{\frac{r'}{\lambda m}} \frac{\ell_*(U_\rho)}{\rho^2} + \frac{r'}{\lambda \rho} \sqrt{\log(e\lambda/r')} \right) \leq \frac{1}{64\lambda},$$

where the last inequality holds as long as

$$\ell_*(U_\rho) \lesssim_L \frac{\rho^2}{\sqrt{r'\lambda}} \sqrt{m} \quad \text{and} \quad r' \sqrt{\log(e\lambda/r')} \lesssim_L \rho. \tag{4.15}$$

It remains to select r' and r'' , taking into account the conditions formulated along the way, specifically, (4.8), (4.9), (4.12), (4.14) and (4.15).

Our starting point is (4.12), which is a condition on ρ and does not involve r' or r'' . Next, the second condition in (4.15) is satisfied by setting

$$r' \simeq_L \frac{\rho}{\sqrt{\log(e\lambda/\rho)}}, \tag{4.16}$$

and one may assume without loss of generality that $r' \leq \rho$ by the choice of $\lambda \geq c(L)$.

With this choice for r' , the first condition in (4.15) holds if

$$\ell_*(U_\rho) \lesssim_L \frac{\rho^{3/2}}{\sqrt{\lambda}} \log^{1/4}(e\lambda/\rho) \cdot \sqrt{m},$$

which is automatically satisfied if (4.12) holds.

With the choice of r' in place, set r'' according to the condition in Lemma 4.6, i.e.,

$$r'' \simeq_L \frac{r'}{\sqrt{\log(e\lambda/r')}} \simeq_L \frac{\rho}{\log(e\lambda/\rho)}. \quad (4.17)$$

Moreover, since $r'' \leq \rho$, (4.9) holds if

$$\ell_*(U_\rho) \lesssim_L \frac{\rho^{3/2}}{\lambda^{1/2} \log^{3/4}(e\lambda/\rho)} \sqrt{m},$$

which is satisfied thanks to (4.12) and the choice $\lambda \geq c(L)$.

Finally, to satisfy (4.8) it suffices that

$$\log |V| = \log \mathcal{N}(T, r'') \lesssim_L \frac{r'}{\lambda} m \simeq_L \frac{m\rho}{\lambda \sqrt{\log(e\lambda/\rho)}},$$

which is true by (4.14). □

The proof of Theorem 1.13 is concluded by combining the estimates on (1), (2) and (3). □

Acknowledgments. The first-named author acknowledges funding from the DFG through the project Quantized Compressive Spectrum Sensing (QuaCoSS) which is part of the Priority Program SPP 1798 Compressive Sensing in Information Processing (COSIP). The authors thank the reviewer for suggestions that helped to improve the presentation of the results.

References

- [1] Ai, A., Lapanowski, A., Plan, Y., Vershynin, R.: One-bit compressed sensing with non-Gaussian measurements. *Linear Algebra Appl.* **441**, 222–239 (2014) [Zbl 1332.94041](#) [MR 3134344](#)
- [2] Artstein-Avidan, S., Giannopoulos, A., Milman, V. D.: *Asymptotic Geometric Analysis. Part I.* *Math. Surveys Monogr.* 202, Amer. Math. Soc., Providence, RI (2015) [Zbl 1337.52001](#) [MR 3331351](#)
- [3] Baraniuk, R. G., Foucart, S., Needell, D., Plan, Y., Wootters, M.: Exponential decay of reconstruction error from binary measurements of sparse signals. *IEEE Trans. Inform. Theory* **63**, 3368–3385 (2017) [MR 3677739](#)
- [4] Boucheron, S., Lugosi, G., Massart, P.: *Concentration Inequalities.* Oxford Univ. Press, Oxford (2013) [Zbl 1279.60005](#) [MR 3185193](#)
- [5] Boufounos, P. T., Baraniuk, R. G.: 1-bit compressive sensing. In: 42nd Annual Conference on Inform. Sciences and Systems, IEEE, 16–21 (2008)

- [6] de la Peña, V. H., Giné, E.: Decoupling. Springer, New York (1999) [Zbl 0918.60021](#) [MR 1666908](#)
- [7] Dirksen, S.: Tail bounds via generic chaining. *Electron. J. Probab.* **20**, no. 53, 29 pp. (2015) [Zbl 1327.60048](#) [MR 3354613](#)
- [8] Dirksen, S., Jung, H. C., Rauhut, H.: One-bit compressed sensing with partial Gaussian circulant matrices. *Inform. Inference* **9**, 601–626 (2020) [MR 4146348](#)
- [9] Dirksen, S., Lecué, G., Rauhut, H.: On the gap between restricted isometry properties and sparse recovery conditions. *IEEE Trans. Inform. Theory* **64**, 5478–5487 (2018) [Zbl 1401.94031](#) [MR 3832317](#)
- [10] Dirksen, S., Mendelson, S.: Robust one-bit compressed sensing with partial circulant matrices. [arXiv:1812.06719](#) (2018)
- [11] Giné, E., Zinn, J.: Some limit theorems for empirical processes. *Ann. Probab.* **12**, 929–998 (1984) [Zbl 0553.60037](#) [MR 757767](#)
- [12] Gray, R. M., Neuhoff, D. L.: Quantization. *IEEE Trans. Inform. Theory* **44**, 2325–2383 (1998) [Zbl 1016.94016](#) [MR 1658787](#)
- [13] Jacques, L., Laska, J. N., Boufounos, P. T., Baraniuk, R. G.: Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Trans. Inform. Theory* **59**, 2082–2102 (2013) [Zbl 1364.94130](#) [MR 3043783](#)
- [14] Knudson, K., Saab, R., Ward, R.: One-bit compressive sensing with norm estimation. *IEEE Trans. Inform. Theory* **62**, 2748–2758 (2016) [Zbl 1359.94117](#) [MR 3493876](#)
- [15] Lecué, G., Mendelson, S.: Sparse recovery under weak moment assumptions. *J. Eur. Math. Soc.* **19**, 881–904 (2017) [Zbl 1414.62135](#) [MR 3612870](#)
- [16] Ledoux, M., Talagrand, M.: Probability in Banach Spaces. *Ergeb. Math. Grenzgeb.* (3) 23, Springer, Berlin (1991) [Zbl 0748.60004](#) [MR 1102015](#)
- [17] Mendelson, S.: Learning without concentration. *J. ACM* **62**, art. 21, 25 pp. (2015) [Zbl 1333.68232](#) [MR 3367000](#)
- [18] Mo, J., Heath, R. W., Jr.: Capacity analysis of one-bit quantized MIMO systems with transmitter channel state information. *IEEE Trans. Signal Process.* **63**, 5498–5512 (2015) [Zbl 1394.94397](#) [MR 3399880](#)
- [19] Oymak, S., Recht, B.: Near-optimal bounds for binary embeddings of arbitrary sets. [CoRR abs/1512.04433](#) (2015)
- [20] Plan, Y., Vershynin, R.: One-bit compressed sensing by linear programming. *Comm. Pure Appl. Math.* **66**, 1275–1297 (2013) [Zbl 1335.94018](#) [MR 3069959](#)
- [21] Plan, Y., Vershynin, R.: Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach. *IEEE Trans. Inform. Theory* **59**, 482–494 (2013) [Zbl 1364.94153](#) [MR 3008160](#)
- [22] Plan, Y., Vershynin, R.: Dimension reduction by random hyperplane tessellations. *Discrete Comput. Geom.* **51**, 438–461 (2014) [Zbl 1296.52014](#) [MR 3164174](#)
- [23] Roberts, L.: Picture coding using pseudo-random noise. *IRE Trans. Inform. Theory* **8**, no. 2, 145–154 (1962)
- [24] Romberg, J.: Compressive sensing by random convolution. *SIAM J. Imaging Sci.* **2**, 1098–1128 (2009) [Zbl 1176.94017](#) [MR 2559161](#)
- [25] Talagrand, M.: Upper and Lower Bounds for Stochastic Processes. *Ergeb. Math. Grenzgeb.* (3) 60, Springer, Heidelberg (2014) [Zbl 1293.60001](#) [MR 3184689](#)