

Sparse recovery under weak moment assumptions

Guillaume Lecué^{1,3} Shahar Mendelson^{2,4,5}

February 23, 2015

Abstract

We prove that iid random vectors that satisfy a rather weak moment assumption can be used as measurement vectors in Compressed Sensing, and the number of measurements required for exact reconstruction is the same as the best possible estimate – exhibited by a random Gaussian matrix. We then show that this moment condition is necessary, up to a log log factor. In addition, we explore the Compatibility Condition and the Restricted Eigenvalue Condition in the noisy setup, as well as properties of neighbourly random polytopes.

1 Introduction and main results

Data acquisition is an important task in diverse fields such as mobile communications, medical imaging, radar detection and others, making the design of efficient data acquisition processes a problem of obvious significance.

The core issue in data acquisition is retaining all the valuable information at one’s disposal, while keeping the ‘acquisition cost’ as low as possible. And while there are several ways of defining that cost, depending on the problem (storage, time, financial cost, etc.), the common denominator of being ‘cost effective’ is ensuring the quality of the data while keeping the number of measurements as small as possible.

The rapidly growing area of *Compressed Sensing* studies ‘economical’ data acquisition processes. We refer the reader to [9, 17] and to the book [23] for more information on the origins of Compressed Sensing and a survey of the progress that has been made in the area in recent years.

¹CNRS, CMAP, Ecole Polytechnique, 91120 Palaiseau, France.

²Department of Mathematics, Technion, I.I.T, Haifa 32000, Israel.

³Email: guillaume.lecue@cmap.polytechnique.fr

⁴Email: shahar@tx.technion.ac.il

⁵Supported by the Mathematical Sciences Institute – The Australian National University and by ISF grant 900/10.

At the heart of Compressed Sensing is a simple idea that has been a recurring theme in Mathematics and Statistics: while complex objects (in this case, data), live in high-dimensional spaces, they can be described effectively using low-dimensional, approximating structures; moreover, randomness may be used to expose these low-dimensional structures. Of course, unlike more theoretical applications of this idea, identifying the low-dimensional structures in the context of Compressed Sensing must be robust and efficient, otherwise, such procedures will be of little practical use.

In the standard Compressed Sensing setup, one observes linear measurements $y_i = \langle X_i, x_0 \rangle$, $i = 1, \dots, N$ of an unknown vector $x_0 \in \mathbb{R}^n$. To make the data acquisition process ‘cost-effective’, the number of measurements N is assumed to be much smaller than the dimension n , and the goal is to identify x_0 using those measurements.

Because the resulting system of equations is under-determined, there is no hope, in general, of identifying x_0 . However, if x_0 is believed to be well approximated by a low-dimensional structure, for example, if x_0 is supported on at most s coordinates for some $s \leq N$, the problem becomes more feasible.

Let (f_1, \dots, f_N) be the canonical basis of \mathbb{R}^N (we will later use (e_1, \dots, e_n) to denote the canonical basis of \mathbb{R}^n) and consider the matrix

$$\Gamma = \frac{1}{\sqrt{N}} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i,$$

called the *measurement matrix*. One possible recovery procedure is ℓ_0 -minimization, in which one selects a vector $t \in \mathbb{R}^n$ that has the shortest support among all vectors satisfying $\Gamma t = \Gamma x_0$. Unfortunately, ℓ_0 minimization is known to be NP-hard in general (see [33] or Theorem 2.17 in [23]). Thus, even without analyzing if and when ℓ_0 -minimization actually recovers x_0 , it is obvious that a more computationally reasonable procedure has to be found.

Fortunately, efficient procedures have been used since the seventies in geophysics (see, for instance [14],[42], [40], and Logan’s Ph.D. thesis [28]). Those procedures are based on ℓ_1 -minimization for which early theoretical works can be found in [20] and [13].

In particular, *Basis Pursuit* is a convex relaxation of ℓ_0 -minimization, and since it can be recast as a linear program (see, e.g., Chapter 15 in [23]), it is far more reasonable than ℓ_0 -minimization from the computational viewpoint.

Definition 1.1 Given the measurement matrix Γ and the measurements $\Gamma x_0 = (\langle X_i, x_0 \rangle)_{i=1}^N$, Basis Pursuit returns a vector \hat{x} that satisfies

$$\hat{x} \in \operatorname{argmin}(\|t\|_1 : \Gamma t = \Gamma x_0). \quad (1.1)$$

Since one may solve this minimization problem effectively, the focus may be shifted to the quality of the solution: whether one can identify measurement vectors X_1, \dots, X_N for which (1.1) has a unique solution, which is x_0 itself, for any x_0 that is s -sparse (i.e. supported on at most s coordinates).

Definition 1.2 Let Σ_s be the set of all s -sparse vectors in \mathbb{R}^n . An $N \times n$ matrix Γ satisfies the **exact reconstruction property of order s** if for every $x_0 \in \Sigma_s$,

$$\operatorname{argmin}(\|t\|_1 : \Gamma t = \Gamma x_0) = \{x_0\}. \quad (\text{ER}(s))$$

It follows from Proposition 2.2.18 in [12] that if Γ satisfies ER(s) then necessarily the number of measurements (rows) is at least $N \geq c_0 s \log(en/s)$, where c_0 is a suitable absolute constant. On the other hand, there are constructions of (random) matrices Γ that satisfy ER(s) with N proportional to $s \log(en/s)$. From here on and with a minor abuse of notation, we will refer to $s \log(en/s)$ as the *optimal number of measurements* and ignore the exact dependence on the constant c_0 .

Unfortunately, the only matrices that are known to satisfy the reconstruction property with an optimal number of measurements are random – which is not surprising, as randomness is one of the most effective tools in exposing low-dimensional, approximating structures. A typical example of an ‘optimal matrix’ is the Gaussian matrix, which has independent standard normal random variables as entries. Other examples of optimal measurement matrices are $\Gamma = N^{-1/2} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i$ where X_1, \dots, X_N are independent, isotropic and L -subgaussian random vectors:

Definition 1.3 A symmetric random vector $X \in \mathbb{R}^n$ is isotropic if for every $t \in \mathbb{R}^n$, $\mathbb{E} \langle X, t \rangle^2 = \|t\|_2^2$; it is L -subgaussian if for every $t \in \mathbb{R}^n$ and every $p \geq 2$, $\|\langle X, t \rangle\|_{L_p} \leq L \sqrt{p} \|\langle X, t \rangle\|_{L_2}$.

The optimal behaviour of isotropic, L -subgaussian matrix ensembles and other ensembles like it, occurs because a typical matrix acts on Σ_s in an isomorphic way when $N \geq c_1 s \log(en/s)$, and in the L -subgaussian case, c_1 is a constant that depends only on L . In Compressed Sensing literature, this isomorphic behaviour is called the *Restricted Isometry property* (RIP) (see, for example [8, 10, 32]): A matrix Γ satisfies the RIP in Σ_s with constant $0 < \delta < 1$, if for every $t \in \Sigma_s$,

$$(1 - \delta) \|t\|_2 \leq \|\Gamma t\|_2 \leq (1 + \delta) \|t\|_2. \quad (1.2)$$

It is straightforward to show that if Γ satisfies the RIP in Σ_{2s} for a sufficiently small constant δ , then it has the exact reconstruction property of order s (see, e.g. [9, 10, 7]).

The standard proof of the RIP for subgaussian ensembles is based on the rapid tail decay of linear functionals $\langle X, t \rangle$. Thus, it seemed natural to ask whether the RIP holds even when linear functionals exhibit a slower decay – for example, when X is L -subexponential – that is, when linear functionals only satisfy that $\|\langle X, t \rangle\|_{L_p} \leq Lp\|\langle X, t \rangle\|_{L_2}$ for every $t \in \mathbb{R}^n$ and every $p \geq 2$.

Proving the RIP for subexponential ensembles is a much harder task than for subgaussian ensembles (see, e.g. [1]). Moreover, the RIP does not exhibit the same optimal quantitative behaviour as in the Gaussian case: it holds with high probability only when $N \geq c_2(L)s \log^2(en/s)$, and this estimate cannot be improved, as can be seen when X has independent, symmetric exponential random variables as coordinates [1].

Although the RIP need not be true for isotropic L -subexponential ensemble using the optimal number of measurements, results in [24] (see Theorem 7.3 there) and in [22] show that exact reconstruction can still be achieved by such an ensemble and *with the optimal number of measurements*. This opens the door to an intriguing question: whether considerably weaker assumptions on the measurement vector may still lead to Exact Reconstruction even when the RIP fails.

The main result presented here does just that, using the *small-ball method* introduced in [30, 31].

Definition 1.4 *A random vector X satisfies the **small-ball condition** in the set Σ_s with constants $u, \beta > 0$ if for every $t \in \Sigma_s$,*

$$P(|\langle X, t \rangle| > u \|t\|_2) \geq \beta.$$

The small-ball condition is a rather minimal assumption on the measurement vector and is satisfied in fairly general situations for values of u and β that are suitable constants, independent of the dimension n .

Under some normalization (like isotropicity), a small-ball condition is an immediate outcome of the Paley-Zygmund inequality (see, e.g. [15]) and moment equivalence. For example, in the following cases a small-ball condition holds with constants that depend only on κ_0 (and on ε for the first case); the straightforward proof may be found in [30].

- X is isotropic and for every $t \in \Sigma_s$, $\|\langle X, t \rangle\|_{L_{2+\varepsilon}} \leq \kappa_0 \|\langle X, t \rangle\|_{L_2}$ for some $\varepsilon > 0$;

- X is isotropic and for every $t \in \Sigma_s$, $\|\langle X, t \rangle\|_{L_2} \leq \kappa_0 \|\langle X, t \rangle\|_{L_1}$.

Because the small-ball condition means that marginals of X do not assign too much weight close to 0, it may hold even without integrability (and in particular, X need not have a covariance matrix). One such example is a random vector whose coordinates are independent random variables that are absolutely continuous with respect to the Lebesgue measure and with a density almost surely bounded by κ_0 . Indeed, as noted in Theorem 1.2 from [38], for every $t \in \mathbb{R}^n$, $\langle X, t/\|t\|_2 \rangle$ has a density that is almost surely bounded by $\sqrt{2}\kappa_0$. In particular, $P(|\langle X, t \rangle| \geq (4\sqrt{2}\kappa_0)^{-1}\|t\|_2) \geq 1/2$ and X satisfies the small ball condition with $u = (4\sqrt{2}\kappa_0)^{-1}$ and $\beta = 1/2$. The estimate on the density of $\langle X, t/\|t\|_2 \rangle$ follows by combining a result due to B. Rogozin [36] on the maximal value of a convolution product of densities, and a result due to K. Ball [2], on the maximal volume of a section of the cube $[-1/2, 1/2]^n$.

Our first result shows that a combination of the small-ball condition and a weak moment assumption suffices to ensure the exact reconstruction property with the optimal number of measurements.

Theorem A. *There exist absolute constants c_0, c_1 and c_2 and for every $\alpha \geq 1/2$ there exists a constant $c_3(\alpha)$ that depends only on α for which the following holds. Let $X = (x_i)_{i=1}^n$ be a random vector on \mathbb{R}^n (with potentially dependent coordinates). Assume that*

1. *there are $\kappa_1, \kappa_2, w > 1$ that satisfy that for every $1 \leq j \leq n$, $\|x_j\|_{L_2} = 1$ and, for every $4 \leq p \leq 2\kappa_2 \log(wn)$, $\|x_j\|_{L_p} \leq \kappa_1 p^\alpha$.*
2. *X satisfies the small ball condition in Σ_s with constants u and β .*

If

$$N \geq c_0 \max \left\{ s \log \left(\frac{en}{s} \right), (c_3(\alpha)\kappa_1^2)^2 (\kappa_2 \log(wn))^{\max\{4\alpha-1, 1\}} \right\},$$

and X_1, \dots, X_N are independent copies of X , then, with probability at least

$$1 - 2 \exp(-c_1 \beta^2 N) - 1/w^{\kappa_2} n^{\kappa_2-1},$$

$\Gamma = N^{-1/2} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i$ satisfies the exact reconstruction property in Σ_{s_1} for $s_1 = c_2 u^2 \beta s$.

An immediate outcome of Theorem A is the following:

- Let x be a centered random variable that has variance 1 and for which $\|x\|_{L_p} \leq c\sqrt{p}$ for $1 \leq p \leq 2 \log n$. If X has independent coordinates distributed as x , then the corresponding matrix Γ with $N \geq c_1 s \log(en/s)$ rows can be used as a measurement matrix and recover any s -sparse vector with large probability.

It is relatively straightforward to derive many other results of a similar flavour, leading to random ensembles that satisfy the exact reconstruction property with the optimal number of measurements.

Remark 1.5 *Our focus is on measurement matrices with independent rows, that satisfy conditions of a stochastic nature – they have i.i.d. rows. Other types of measurement matrices that have some structure have also been used in Compressed Sensing. One notable example is a random Fourier measurement matrix, obtained by randomly selecting rows from the discrete Fourier matrix (see, e.g. [10], [37] or Chapter 12 in [23]).*

One may wonder if the small-ball condition is satisfied for more structured matrices, as the argument we use here does not extend immediately to such cases. And, indeed, for structured ensembles one may encounter a different situation: a small-ball condition that is not uniform, in the sense that the constants u and β from Definition 1.4 are direction-dependent. Moreover, in some cases, the known estimates on these constants are far from what is expected.

Results of the same flavour of Theorem A may follow from a ‘good enough’ small-ball condition, even if it is not uniform, by slightly modifying the argument we use here. However, obtaining a satisfactory ‘non-uniform’ small-ball condition is a different story. For example, in the Fourier case, such an estimate is likely to require quantitative extensions of the Littlewood-Paley theory – a worthy challenge in its own right, and one which goes far beyond the goals of this article.

Just as noted for subexponential ensembles, Theorem A cannot be proved using an RIP-based argument. A key ingredient in the proof is the following observation:

Theorem B. *Let $\Gamma : \mathbb{R}^n \mapsto \mathbb{R}^N$ and denote by (e_1, \dots, e_n) the canonical basis of \mathbb{R}^n . Assume that:*

- for every $x \in \Sigma_s$, $\|\Gamma x\|_2 \geq c_0 \|x\|_2$, and*
- for every $j \in \{1, \dots, n\}$, $\|\Gamma e_j\|_2 \leq c_1$.*

Setting $s_1 = \lfloor (c_0^2(s-1))/(4c_1^2) \rfloor - 1$, Γ satisfies the exact reconstruction property in Σ_{s_1} .

Compared with the RIP, conditions a) and b) in Theorem B are weaker, as it suffices to verify the right-hand side of (1.2) for 1-sparse vectors rather than for every s -sparse vector. This happens to be a substantial difference: the assumption that for every $t \in \Sigma_s$, $\|\Gamma t\|_2 \leq (1 + \delta) \|t\|_2$ is a costly one, and happens to be the reason for the gap between the RIP and the exact reconstruction property. Indeed, while the lower bound in the RIP holds for rather general ensembles (see [30] and the next section for more details), and is guaranteed solely by the small-ball condition, the upper bound is almost equivalent to having the coordinates of X exhibit a subgaussian behaviour of moments, at least up to some level. Even the fact that one has to verify the upper bound for 1-sparse vectors comes at a cost, namely, the moment assumption (1) in Theorem A.

The second goal of this note is to illustrate that while Exact Reconstruction is ‘cheaper’ than the RIP, it still comes at a cost – namely, that the moment condition (1) in Theorem A is truly needed.

Definition 1.6 *A random matrix Γ is generated by the random variable x if $\Gamma = N^{-1/2} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i$ and X_1, \dots, X_N are independent copies of the random vector $X = (x_1, \dots, x_n)^\top$ whose coordinates are independent copies of x .*

Theorem C. *There exist absolute constants c_0, c_1, c_2 and c_3 for which the following holds. Given $n \geq c_0$ and $N \log N \leq c_1 n$, there exists a mean-zero, variance one random variable x with the following properties:*

- $\|x\|_{L_p} \leq c_2 \sqrt{p}$ for $2 < p \leq c_3(\log n)/(\log N)$.
- If $(x_j)_{j=1}^n$ are independent copies of x then $X = (x_1, \dots, x_n)^\top$ satisfies the small-ball condition with constants u and β that depend only on c_2 .
- Denote by Γ the $N \times n$ matrix generated by x . For every $k \in \{1, \dots, n\}$, with probability larger than $1/2$, $\operatorname{argmin}(\|t\|_1 : \Gamma t = \Gamma e_k) \neq \{e_k\}$; therefore, e_k is not exactly reconstructed by Basis Pursuit and so Γ does not satisfy the exact reconstruction property of order 1.

To put Theorem C in some perspective, note that if Γ is generated by x for which $\|x\|_{L_2} = 1$ and $\|x\|_{L_p} \leq c_4 \sqrt{p}$ for $2 < p \leq c_5 \log n$, then $X =$

$(x_i)_{i=1}^n$ satisfies the small-ball condition with constants that depend only on c_4 , and by Theorem A, if $N \geq c_6 \log n$, Γ satisfies ER(1) with high probability. On the other hand, the random ensemble from Theorem C is generated by x that has almost identical properties – with one exception: its L_p norm is well behaved only for $p \leq c_7(\log n)/\log \log n$. This small gap in the number of moments has a significant impact: with probability at least $1/2$, Γ does not satisfy ER(1) when N is of the order of $\log n$.

Therefore, the moment condition in Theorem A is indeed required (up to a $\log \log n$ factor).

The idea behind the proof of Theorem C is to construct a random matrix Γ for which, given any basis vector e_k , with probability at least $1/2$, $\|\Gamma e_k\|_2 \leq 1$, while the set $\{\Gamma e_j, j \neq k\}$ has many ‘very spiky’ vectors: the convex hull $\text{conv}(\pm \Gamma e_j, j \neq k)$ contains a perturbation of $2\sqrt{N}B_1^N$, i.e., a large multiple of the unit ball in ℓ_1^N . Since such a set must contain the Euclidean unit ball, and in particular, Γe_k as well, it follows that e_k cannot be the unique solution of the ℓ_1 minimization problem $\min(\|t\|_1 : \Gamma t = \Gamma e_k)$.

The fact that the coordinates of X do not have enough well behaved moments is the key feature that allows one to generate many ‘spiky’ columns in a typical Γ .

An alternative formulation of Theorem C is the following:

Theorem C’. *There are absolute constants c_0, c_1, c_2 and κ for which the following holds. If $n \geq c_0$ and $2 < p < c_1 \log n$, there exists a mean-zero and variance 1 random variable x , for which $\|x\|_{L_q} \leq \kappa\sqrt{q}$ for $2 < q \leq p$, and if $N \leq c_2\sqrt{p}(n/\log n)^{1/p}$ and Γ is the $N \times n$ matrix generated by x , then with probability at least $1/2$, Γ does not satisfy the exact reconstruction property of order 1.*

Theorem C and Theorem C’ imply that Basis Pursuit may perform poorly when the coordinates of X do not have enough moments, and requires a polynomial number of measurements in n to ensure Exact Reconstruction. This happens to be the price of convex relaxation: a rather striking observation is that ℓ_0 -minimization achieves recovery with the optimal number of measurements under an even weaker small-ball condition than in Theorem A, and without any additional moment assumptions.

Recall that ℓ_0 -minimization is defined by $\hat{x} = \text{argmin}(\|t\|_0 : \Gamma t = \Gamma x_0)$, where $\|t\|_0$ is the cardinality of the support of t .

Definition 1.7 *X satisfies a weak small-ball condition in Σ_s with constant*

β if for every $t \in \Sigma_s$,

$$P(|\langle X, t \rangle| > 0) \geq \beta. \quad (1.3)$$

Theorem D. *For every $0 < \beta < 1$ there exist constants c_0 and c_1 that depend only on β and for which the following holds. Let X be a random vector that satisfies the weak small-ball condition in Σ_s with a constant β . Let X_1, \dots, X_N be N independent copies of X and set $\Gamma = N^{-1/2} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i$. If $N \geq c_0 s \log(en/s)$ then with probability at least $1 - 2 \exp(-c_1 N)$, for every $x_0 \in \Sigma_{\lfloor s/2 \rfloor}$, ℓ_0 -minimization has a unique solution, which is x_0 itself.*

The price of convex relaxation can now be clearly seen through the number of measurements needed for exact reconstruction: consider the random vector X constructed in Theorem C' for, say, $p = 4$. Since X satisfies the conditions of Theorem D, ℓ_0 minimization may be used to recover any s -sparse vector with only $N = cs \log(en/s)$ random measurements. In contrast, Basis Pursuit requires at least $\sim (n/\log n)^{1/4}$ measurements to reconstruct 1-sparse vectors.

It should be noted that under much stronger assumptions on X , the exact recovery of s -sparse vectors using ℓ_0 -minimization may occur when N is as small as $2s$. Indeed, it suffices to ensure that all the $N \times 2s$ sub-matrices of Γ are non-singular, and this is the case when $N = 2s$ if the entries of Γ are independent random variables that are absolutely continuous (see Chapter 2 in [23] for more details).

We end this introduction with a word about notation and the organization of the article. The proofs of Theorem A, Theorem B and Theorem D are presented in the next section, while the proofs of Theorem C and Theorem C' may be found in Section 3. The final section is devoted to results in a natural 'noisy' extension of Compressed Sensing. In particular, we prove that both the *Compatibility Condition* and the *Restricted Eigenvalue Condition* hold under weak moment assumptions; we also study related properties of random polytopes.

As for notation, throughout, absolute constants or constants that depend on other parameters are denoted by c, C, c_1, c_2 , etc., (and, of course, we will specify when a constant is absolute and when it depends on other parameters). The values of these constants may change from line to line. The notation $x \sim y$ (resp. $x \lesssim y$) means that there exist absolute constants $0 < c < C$ for which $cy \leq x \leq Cy$ (resp. $x \leq Cy$). If $b > 0$ is a parameter then $x \lesssim_b y$ means that $x \leq C(b)y$ for some constant $C(b)$ that depends only on b .

Let ℓ_p^m be \mathbb{R}^m endowed with the norm $\|x\|_{\ell_p^m} = (\sum_j |x_j|^p)^{1/p}$; the corresponding unit ball is denoted by B_p^m and the unit Euclidean sphere in \mathbb{R}^m is S^{m-1} . If $A \subset \mathbb{R}^n$ then $\mathbb{1}_A$ denotes the indicator function of A . Finally, we will assume that (\mathcal{X}, μ) is a probability space, and that X is distributed according to μ .

2 Proof of Theorem A, B and D

The proof of Theorem A has several components, and although the first of which is rather standard, we present it for the sake of completeness.

Lemma 2.1 *Let $\Gamma : \mathbb{R}^n \rightarrow \mathbb{R}^N$ be a matrix and set $\ker(\Gamma)$ to be its kernel. If $0 < r < 1$ and $B_1^n \cap rS^{n-1}$ does not intersect $\ker(\Gamma)$, then Γ satisfies the exact reconstruction property in $\Sigma_{\lfloor (2r)^{-2} \rfloor}$.*

Proof. Observe that if $x \in B_1^n$ and $\|x\|_2 \geq r$ then $y = rx/\|x\|_2 \in B_1^n \cap rS^{n-1}$. Therefore, if $y \notin \ker(\Gamma)$, the same holds for x ; thus

$$\sup_{x \in B_1^n \cap \ker(\Gamma)} \|x\|_2 < r.$$

Let $s = \lfloor (2r)^{-2} \rfloor$, fix $x_0 \in \Sigma_s$ and put I to be the set indices of coordinates on which x_0 is supported. Given a nonzero $h \in \ker(\Gamma)$, let $h = h_I + h_{I^c}$ – the decomposition of h to coordinates in I and in I^c . Since $h/\|h\|_1 \in B_1^n \cap \ker(\Gamma)$, it follows that $\|h\|_2 < r\|h\|_1$, and by the choice of s , $2\sqrt{s}\|h\|_2 < \|h\|_1$. Therefore,

$$\begin{aligned} \|x_0 + h\|_1 &= \|x_0 + h_I\|_1 + \|h_{I^c}\|_1 \geq \|x_0\|_1 - \|h_I\|_1 + \|h_{I^c}\|_1 \\ &= \|x_0\|_1 - 2\|h_I\|_1 + \|h\|_1 \geq \|x_0\|_1 - 2\sqrt{|I|}\|h\|_2 + \|h\|_1 > \|x_0\|_1. \end{aligned}$$

Hence, $\|x_0 + h\|_1 > \|x_0\|_1$ and x_0 is the unique minimizer of the basis pursuit algorithm. \blacksquare

The main ingredient in the proof of Theorem A is Lemma 2.3 below, which is based on the small-ball method introduced in [30, 31]. To formulate the lemma, one requires the notion of a VC class of sets.

Definition 2.2 *Let \mathcal{G} be a class of $\{0, 1\}$ -valued functions defined on a set \mathcal{X} . The set \mathcal{G} is a VC-class if there exists an integer V for which, given any $x_1, \dots, x_{V+1} \in \mathcal{X}$,*

$$|\{(g(x_1), \dots, g(x_{V+1})) : g \in \mathcal{G}\}| < 2^{V+1}. \quad (2.1)$$

The VC-dimension of \mathcal{G} , denoted by $VC(\mathcal{G})$, is the smallest integer V for which (2.1) holds.

The VC dimension is a combinatorial complexity measure that may be used to control the $L_2(\mu)$ -covering numbers of the class; indeed, set $N(\mathcal{G}, \varepsilon, L_2(\mu))$ to be the smallest number of open balls of radius ε relative to the $L_2(\mu)$ norm that are needed to cover \mathcal{G} . A well known result due to Dudley [21] is that if $VC(\mathcal{G}) = V$ and μ is a probability measure on \mathcal{X} then for every $0 < \varepsilon < 1$,

$$N(\mathcal{G}, \varepsilon, L_2(\mu)) \leq \left(\frac{c_1}{\varepsilon}\right)^{c_2 V}, \quad (2.2)$$

where c_1 and c_2 are absolute constants.

Lemma 2.3 *There exist absolute constants c_1 and c_2 for which the following holds. Let \mathcal{F} be a class of functions and assume that there are $\beta > 0$ and $u \geq 0$ for which*

$$\inf_{f \in \mathcal{F}} P(|f(X)| > u) \geq \beta.$$

Let $\mathcal{G}_u = \{\mathbb{1}_{\{|f| > u\}} : f \in \mathcal{F}\}$. If $VC(\mathcal{G}_u) \leq d$ and $N \geq c_1 d / \beta^2$ then with probability at least $1 - \exp(-c_2 \beta^2 N)$,

$$\inf_{f \in \mathcal{F}} |\{i \in \{1, \dots, N\} : |f(X_i)| > u\}| \geq \frac{\beta N}{2}.$$

Remark 2.4 *Note that $u = 0$ is a ‘legal choice’ in Lemma 2.3, a fact that will be used in the proof of Theorem D.*

Proof. Let $G(X_1, \dots, X_N) = \sup_{g \in \mathcal{G}_u} |N^{-1} \sum_{i=1}^N g(X_i) - \mathbb{E}g(X)|$. By the bounded differences inequality (see, for example, Theorem 6.2 in [5]), with probability at least $1 - \exp(-t)$,

$$G(X_1, \dots, X_N) \leq \mathbb{E}G(X_1, \dots, X_N) + c_1 \sqrt{\frac{t}{N}}.$$

Standard empirical processes arguments (symmetrization, the fact that Bernoulli processes are subgaussian and the entropy estimate (2.2) – see, for example, Chapters 2.2, 2.3 and 2.6 in [47]), show that since $VC(\mathcal{G}) \leq d$,

$$\mathbb{E}G(X_1, \dots, X_N) \leq c_2 \sqrt{\frac{d}{N}} \leq \frac{\beta}{4}, \quad (2.3)$$

provided that $N \gtrsim d / \beta^2$. Therefore, taking $t = N\beta^2 / 16c_1^2$, it follows that with probability at least $1 - \exp(-c_3 \beta^2 N)$, for every $f \in \mathcal{F}$,

$$\frac{1}{N} \sum_{i=1}^N \mathbb{1}_{\{|f| > u\}}(X_i) \geq P(|f(X)| > u) - \frac{\beta}{2} \geq \frac{\beta}{2}.$$

Therefore, on that event, $|\{i : |f(X_i)| > u\}| \geq \beta N / 2$ for every $f \in \mathcal{F}$. ■

Corollary 2.5 *There exist absolute constants c_1 and c_2 for which the following holds. Let $X \in \mathbb{R}^n$ be a random vector.*

1. *If there are $0 < \beta \leq 1$ and $u \geq 0$ for which $P(|\langle t, X \rangle| > u) \geq \beta$ for every $t \in S^{n-1}$ and if $N \geq c_1 n / \beta^2$, then with probability at least $1 - \exp(-c_2 N \beta^2)$,*

$$\inf_{t \in S^{n-1}} \frac{1}{N} \sum_{i=1}^N \langle X_i, t \rangle^2 > \frac{u^2 \beta}{2}.$$

2. *If there are $0 < \beta \leq 1$ and $u \geq 0$ for which $P(|\langle t, X \rangle| > u) \geq \beta$ for every $t \in \Sigma_s \cap S^{n-1}$ and if $N \geq c_1 s \log(en/s) / \beta^2$, then with probability at least $1 - \exp(-c_2 N \beta^2)$,*

$$\inf_{t \in \Sigma_s \cap S^{n-1}} \frac{1}{N} \sum_{i=1}^N \langle X_i, t \rangle^2 > \frac{u^2 \beta}{2}.$$

Remark 2.6 *Note that the first part of Corollary 2.5 gives an estimate on the smallest singular value of the random matrix $\Gamma = N^{-1/2} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i$. The proof follows the same path as in [25], but unlike the latter, no assumption on the covariance structure of X , used both in [25] and in [41], is required. In fact, Corollary 2.5 may be applied even if the covariance matrix does not exist. Thus, under a small-ball condition, the smallest singular value of Γ is larger than $c(\beta, u)$ with high (exponential) probability.*

Proof of Corollary 2.5. To prove the first part of the claim, let $\mathcal{F} = \{\langle t, \cdot \rangle : t \in S^{n-1}\}$. Recall that the VC dimension of a class of half-spaces in \mathbb{R}^n is at most n , and thus, one may verify that for every $u \geq 0$, the VC dimension of

$$\mathcal{G}_u = \{\mathbb{1}_{\{|f| > u\}} : f \in \mathcal{F}\}$$

is at most $c_1 n$ for a suitable absolute constant c_1 (see, e.g., Chapter 2.6 in [47]). The claim now follows immediately from Lemma 2.3 because

$$\frac{1}{N} \sum_{i=1}^N \langle t, X_i \rangle^2 > \frac{u^2}{N} |\{i : |\langle X_i, t \rangle| > u\}|$$

for every $t \in S^{n-1}$.

Turning to the second part, note that $\Sigma_s \cap S^{n-1}$ is a union of $\binom{n}{s}$ spheres of dimension s . Applying the first part to each one of those spheres, combined with the union bound, it follows that for $N \geq c_2 \beta^{-2} s \log(en/s)$, with probability at least $1 - \exp(-c_3 N \beta^2)$,

$$\inf_{t \in \Sigma_s \cap S^{n-1}} \frac{1}{N} \sum_{i=1}^N \langle X_i, t \rangle^2 > \frac{u^2 \beta}{2}.$$

■

Corollary 2.5 shows that the small-ball condition for linear functionals implies that Γ ‘acts well’ on s -sparse vectors. However, according to Lemma 2.1, exact recovery is possible if Γ is well behaved on the set

$$\sqrt{\kappa_0 s} B_1^n \cap S^{n-1} = \{x \in \mathbb{R}^n : \|x\|_1 \leq \sqrt{\kappa_0 s}, \|x\|_2 = 1\}$$

for a well-chosen constant κ_0 . In the standard (RIP-based) argument, one proves exact reconstruction by first showing that the RIP holds in Σ_s , and then the fact that each vector in $\sqrt{\kappa_0 s} B_1^n \cap S^{n-1}$ is well approximated by vectors from Σ_s (see, for instance, [12]) allows one to extend the RIP from Σ_s to $\sqrt{\kappa_0 s} B_1^n \cap S^{n-1}$. Unfortunately, this extension requires *both upper and lower* estimates in the RIP.

Since the upper part of the RIP in Σ_s forces severe restrictions on the random vector X , one has to resort to a totally different argument if one wishes to extend the *lower bound* from Σ_s (which only requires the small-ball condition) to $\sqrt{\kappa_0 s} B_1^n \cap S^{n-1}$.

The method presented below is based on Maurey’s empirical method and has been recently used in [34].

Lemma 2.7 *Let $\Gamma : \mathbb{R}^n \rightarrow \mathbb{R}^N$, put $1 < s \leq n$ and assume that for every $x \in \Sigma_s$, $\|\Gamma x\|_2 \geq \lambda \|x\|_2$. If $y \in \mathbb{R}^n$ is a nonzero vector and $\mu_j = |y_j|/\|y\|_1$, then*

$$\|\Gamma y\|_2^2 \geq \lambda^2 \|y\|_2^2 - \frac{\|y\|_1^2}{s-1} \left(\sum_{j=1}^n \|\Gamma e_j\|_2^2 \mu_j - \lambda^2 \right).$$

Proof. Fix $y \in \mathbb{R}^n$, let Y be a random vector in \mathbb{R}^n defined by

$$P(Y = \|y\|_1 \operatorname{sgn}(y_j) e_j) = |y_j|/\|y\|_1,$$

for every $j = 1, \dots, n$ and observe that $\mathbb{E}Y = y$.

Let Y_1, \dots, Y_s be independent copies of Y and set $Z = s^{-1} \sum_{k=1}^s Y_k$. Note that $Z \in \Sigma_s$ for every realization of Y_1, \dots, Y_s ; thus $\|\Gamma Z\|_2^2 \geq \lambda^2 \|Z\|_2^2$ and

$$\mathbb{E}\|\Gamma Z\|_2^2 \geq \lambda^2 \mathbb{E}\|Z\|_2^2. \quad (2.4)$$

It is straightforward to verify that $\mathbb{E}\langle Y, Y \rangle = \|y\|_1^2$; that if $i \neq j$ then $\mathbb{E}\langle \Gamma Y_i, \Gamma Y_j \rangle = \langle \Gamma y, \Gamma y \rangle$; and that for every $1 \leq k \leq s$,

$$\mathbb{E}\langle \Gamma Y_k, \Gamma Y_k \rangle = \|y\|_1 \sum_{j=1}^n |y_j| \|\Gamma e_j\|_2^2.$$

Therefore, setting $\mu_j = |y_j|/\|y\|_1$ and $W = \sum_{j=1}^n \|\Gamma e_j\|_2^2 \mu_j$,

$$\begin{aligned} \mathbb{E}\|\Gamma Z\|_2^2 &= \frac{1}{s^2} \sum_{i,j=1}^s \mathbb{E}\langle \Gamma Y_i, \Gamma Y_j \rangle = \left(1 - \frac{1}{s}\right) \|\Gamma y\|_2^2 + \frac{\|y\|_1}{s} \sum_{j=1}^n |y_j| \|\Gamma e_j\|_2^2 \\ &= \left(1 - \frac{1}{s}\right) \|\Gamma y\|_2^2 + W \frac{\|y\|_1^2}{s}, \end{aligned}$$

and using the same argument one may show that

$$\mathbb{E}\|Z\|_2^2 = \left(1 - \frac{1}{s}\right) \|y\|_2^2 + \frac{\|y\|_1^2}{s}.$$

Combining these two estimates with (2.4),

$$\left(1 - \frac{1}{s}\right) \|\Gamma y\|_2^2 \geq \lambda^2 \left(\left(1 - \frac{1}{s}\right) \|y\|_2^2 + \frac{\|y\|_1^2}{s} \right) - W \frac{\|y\|_1^2}{s},$$

proving the claim. \blacksquare

Proof of Theorem B: Assume that for every $x \in \Sigma_s$, $\|\Gamma x\|_2 \geq c_0 \|x\|_2$ and that for every $1 \leq i \leq n$, $\|\Gamma e_i\|_2 \leq c_1$. It follows from Lemma 2.7 that if $s - 1 > c_1^2/(c_0^2 r^2)$, then for every $y \in B_1^n \cap rS^{n-1}$,

$$\|\Gamma y\|_2^2 \geq c_0^2 \|y\|_2^2 - \frac{\|y\|_1}{s-1} \sum_{i=1}^n \|\Gamma e_i\|_2^2 |y_i| \geq c_0^2 r^2 - \frac{c_1^2}{s-1} > 0.$$

The claim now follows from Lemma 2.1. \blacksquare

Consider the matrix $\Gamma = N^{-1/2} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i$. Observe that for every $t \in \mathbb{R}^n$, $\|\Gamma t\|_2^2 = N^{-1} \sum_{i=1}^N \langle X_i, t \rangle^2$, and that if $X_j = (x_{i,j})_{i=1}^n$ then

$$\|\Gamma e_j\|_2^2 = \frac{1}{N} \sum_{i=1}^N x_{i,j}^2,$$

which is an average of N iid random variables (though $\|\Gamma e_1\|_2, \dots, \|\Gamma e_n\|_2$ need not be independent).

Thanks to Theorem B and Corollary 2.5, the final component needed for the proof of Theorem A is information on the sum of iid random variables, which will be used to bound $\max_{1 \leq j \leq n} \|\Gamma e_j\|_2^2$ from above.

Lemma 2.8 *There exists an absolute constant c_0 for which the following holds. Let z be a mean-zero random variable and put z_1, \dots, z_N to be N independent copies of z . Let $p_0 \geq 2$ and assume that there exists $\kappa_1 > 0$ and $\alpha \geq 1/2$ for which $\|z\|_{L_p} \leq \kappa_1 p^\alpha$ for every $2 \leq p \leq p_0$. If $N \geq p_0^{\max\{2\alpha-1, 1\}}$ then for every $2 \leq p \leq p_0$,*

$$\left\| \frac{1}{\sqrt{N}} \sum_{i=1}^N z_i \right\|_{L_p} \leq c_1(\alpha) \kappa_1 \sqrt{p},$$

where $c_1(\alpha) = c_0 \exp((2\alpha - 1))$.

Lemma 2.8 shows that even under a weak moment assumption, namely that $\|z\|_{L_p} \lesssim p^\alpha$ for $p \leq p_0$ and $\alpha \geq 1/2$ that can be large, a normalized sum of N independent copies of z exhibits a ‘subgaussian’ moment growth up to the same p_0 , as long as N is sufficiently large.

The proof of Proposition 2.8 is based on the following fact due to Latała.

Theorem 2.9 ([26], Theorem 2 and Remark 2) *If z is a mean-zero random variable and z_1, \dots, z_N are independent copies of z , then for any $p \geq 2$,*

$$\left\| \sum_{i=1}^N z_i \right\|_{L_p} \sim \sup \left\{ \frac{p}{s} \left(\frac{N}{p} \right)^{1/s} \|z\|_{L_s} : \max\{2, p/N\} \leq s \leq p \right\}.$$

Proof of Lemma 2.8. Let $2 \leq p \leq p_0$ and $N \geq p$. Since $\|z\|_{L_s} \leq \kappa_1 s^\alpha$ for any $2 \leq s \leq p$, it follows from Theorem 2.9 that

$$\left\| \sum_{i=1}^N z_i \right\|_{L_p} \leq c_0 \kappa_1 \sup \left\{ p(N/p)^{1/s} s^{-1+\alpha} : \max\{2, p/N\} \leq s \leq p \right\}.$$

It is straightforward to verify that the function $h(s) = (N/p)^{1/s} s^{-1+\alpha}$ is non-increasing when $\alpha \leq 1$ and attains its maximum in $s = \max\{2, p/N\} = 2$ or in $s = p$ when $\alpha > 1$. Therefore, when $N \geq p$,

$$\left\| \sum_{i=1}^N z_i \right\|_{L_p} \leq c_1 \kappa_1 \max \left\{ \sqrt{Np}, N^{1/p} p^\alpha \right\}.$$

Finally, if $N \geq p^{2\alpha-1}$ then $e^{2\alpha-1} \sqrt{Np} \geq N^{1/p} p^\alpha$, which completes the proof. ■

Proof of Theorem A. Consider $N \geq c_1 s \log(en/s)/\beta^2$. By Corollary 2.5, with probability at least $1 - \exp(-c_2 N \beta^2)$,

$$\inf_{t \in \Sigma_s \cap S^{n-1}} \frac{1}{N} \sum_{i=1}^N \langle X_i, t \rangle^2 > \frac{u^2 \beta}{2}. \quad (2.5)$$

Set $(X_i)_{i=1}^N$ for which (2.5) holds and let $\Gamma = N^{-1/2} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i$. By Lemma 2.7 for $\lambda^2 = u^2 \beta / 2$, it follows that when $r \geq 1$,

$$\inf_{t \in \sqrt{r} B_1^n \cap S^{n-1}} \|\Gamma t\|_2^2 \geq \lambda^2 - \frac{2r}{s} \max_{1 \leq j \leq n} \|\Gamma e_j\|_2^2. \quad (2.6)$$

Next, one has to obtain a high probability upper estimate on $\max_{1 \leq j \leq n} \|\Gamma e_j\|_2^2$. To that end, fix $w \geq 1$ and consider $z = x_j^2 - 1$ - where x_j is the j -th coordinate of X . Observe that z is a centered random variable and that $\|z\|_{L_q} \lesssim 4^\alpha \kappa_1^2 q^{2\alpha}$ for every $1 \leq q \leq \kappa_2 \log(wn)$. Thus, by Lemma 2.8 for $p = \kappa_2 \log(wn)$ and $c_3(\alpha) \sim 4^\alpha \exp((4\alpha - 1))$,

$$\left\| \frac{1}{N} \sum_{i=1}^N z_i \right\|_{L_p} \leq c_3(\alpha) \kappa_1^2 \sqrt{\frac{p}{N}},$$

provided that $N \geq p^{\max\{4\alpha-1, 1\}} = (\kappa_2 \log(wn))^{\max\{4\alpha-1, 1\}}$. Hence, if $N \geq (c_3(\alpha) \kappa_1^2)^2 (\kappa_2 \log(wn))^{\max\{4\alpha-1, 1\}}$, and setting $V_j = \|\Gamma e_j\|_2^2$, one has

$$\|V_j\|_{L_p} = \left\| \frac{1}{N} \sum_{i=1}^N x_{i,j}^2 \right\|_{L_p} \leq 1 + c_3(\alpha) \kappa_1^2 \sqrt{\frac{\kappa_2 \log(wn)}{N}} \leq 2;$$

thus,

$$\begin{aligned} P\left(\max_{1 \leq j \leq n} V_j \geq 2e\right) &\leq \sum_{j=1}^n P(V_j \geq 2e) \leq \sum_{j=1}^n \left(\frac{\|V_j\|_{L_p}}{2e}\right)^p \\ &\leq n \left(\frac{1}{e}\right)^p = \frac{1}{w^{\kappa_2 n^{\kappa_2-1}}}. \end{aligned}$$

Combining the two estimates, if

$$N \gtrsim \max \left\{ s \log(en/s), (c_3(\alpha) \kappa_1^2)^2 (\kappa_2 \log(wn))^{\max\{4\alpha-1, 1\}} \right\}$$

and $r \leq s \lambda^2 / 8e = s u^2 \beta / 16e$, then with probability at least $1 - \exp(-c_2 N \beta^2) - 1/(w^{\kappa_2 n^{\kappa_2-1}})$,

$$\inf_{t \in \sqrt{r} B_1^n \cap S^{n-1}} \|\Gamma t\|_2^2 \geq \lambda^2 - \frac{4er}{s} \geq \lambda^2 / 2. \quad (2.7)$$

Therefore, by Lemma 2.1, Γ satisfies the exact reconstruction property for vectors that are $c_4 u^2 \beta s$ -sparse, as claimed. \blacksquare

Proof of Theorem D: Since the argument is almost identical to the one used in the proof of the second part of Corollary 2.5, we will only sketch the details. Observe that if $\Gamma = N^{-1/2} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i$ and $\ker(\Gamma) \cap \Sigma_s = \{0\}$, then for any $x_0 \in \Sigma_{\lfloor s/2 \rfloor}$, the only $t \in \mathbb{R}^n$ for which $\Gamma t = \Gamma x_0$ and $\|t\|_0 \leq \|x_0\|_0$ is x_0 itself. Thus, it suffices to show that for every $x \in \Sigma_s \cap S^{n-1}$, $|\langle X_i, x \rangle| > 0$ for some $1 \leq i \leq n$. Since $\Sigma_s \cap S^{n-1}$ is a union of $\binom{n}{s}$ spheres, the claim follows from Lemma 2.3 applied to each one of those spheres and for $u = 0$, combined with a union bound argument. ■

3 Proof of Theorem C and Theorem C'

Consider an $N \times n$ matrix Γ and $J \subset \{1, \dots, n\}$. Set Γ_J to be the $(N \times |J|)$ restriction of Γ to $\text{span}\{e_j : j \in J\}$. Recall that B_1^n is the unit ball in $\ell_1^n = (\mathbb{R}^n, \|\cdot\|_1)$, and put $B_1^{J^c}$ to be the set of vectors in B_1^n that are supported in J^c – the complement of J in $\{1, \dots, n\}$.

Lemma 3.1 *Fix integers $s, N \leq n$ and $J \subset \{1, \dots, n\}$ of cardinality at most s . If $v \in \mathbb{R}^n$ is supported in J , $\|v\|_1 = 1$ and $\Gamma v \in \Gamma B_1^{J^c}$, then Γ does not satisfy the exact reconstruction property of order s .*

Proof. Let $w \in B_1^{J^c}$ for which $\Gamma v = \Gamma w$ and observe that $v \neq w$ (otherwise, $v \in B_1^J \cap B_1^{J^c}$, implying that $v = 0$, which is impossible because $\|v\|_1 = 1$).

Since $\|w\|_1 \leq 1 = \|v\|_1$, w is at least as good a candidate as v for the ℓ_1 -minimization problem $\min(\|t\|_1 : \Gamma t = \Gamma v)$; hence, v is not the unique solution of that problem. ■

Set x_1, \dots, x_n to be the columns of Γ . It immediately follows from Lemma 3.1 that if one wishes to prove that Γ does not satisfy ER(1), it suffices to show that, for instance the first basis vector e_1 cannot be exactly reconstruct. This follows from

$$\Gamma e_1 = x_1 \in \text{absconv}(x_k : k \neq 1) = \text{absconv}(\Gamma e_k : k \neq 1) = \Gamma B_1^{\{1\}^c},$$

where $\text{absconv}(S)$ is the convex hull of $S \cup -S$. Therefore, if

$$\|x_1\|_2 \leq c_0 \quad \text{and} \quad c_0 B_2^N \subset \text{absconv}(x_k : k \neq 1), \quad (3.1)$$

for some absolute constant c_0 , then Γ does not satisfy ER(1).

The proofs of Theorem C and of Theorem C' follow from the construction of a random matrix ensemble for which (3.1) holds with probability larger than $1/2$. We now turn on to such a construction.

Let η be a selector (a $\{0, 1\}$ -valued random variable) with mean δ to be named later, and let ε be a symmetric $\{-1, 1\}$ -valued random variable that is independent of η . Fix $R > 0$ and set

$$z = \varepsilon(1 + R\eta).$$

Observe that if $p \geq 2$ and $R \geq 1$ then

$$\frac{\|z\|_{L_p}}{\|z\|_{L_2}} = \frac{(1 + ((1 + R)^p - 1)\delta)^{1/p}}{(1 + ((1 + R)^2 - 1)\delta)^{1/2}} \sim \frac{(1 + R^p\delta)^{1/p}}{(1 + R^2\delta)^{1/2}} \sim R\delta^{1/p},$$

and the last equivalence holds when $R^2\delta \lesssim 1$ and $R^p\delta \gtrsim 1$. Fix $2 < p \leq 2 \log(1/\delta)$ which will be specified later and set $R = \sqrt{p}(1/\delta)^{1/p}$. Since the function $q \rightarrow \sqrt{q}/\delta^{1/q}$ is decreasing for $2 \leq q \leq 2 \log(1/\delta)$ one has that for $2 \leq q \leq p$ and for δ that is small enough,

$$\|z\|_{L_q} \leq c_0 \sqrt{q} \|z\|_{L_2}.$$

Note that $x = z/\|z\|_{L_2}$ is a mean-zero, variance one random variable that exhibits a ‘subgaussian’ moment behaviour only up to p . Indeed, if $2 \leq q \leq p$, $\|z\|_{L_q} \lesssim \sqrt{q} \|z\|_{L_2}$, and if $q > p$, $\|z\|_{L_q} \sim \sqrt{p} \delta^{1/q-1/p} \|z\|_{L_2}$, which may be far larger than $\sqrt{q} \|z\|_{L_2}$ if δ is sufficiently small.

Let $X = (x_1, \dots, x_n)$ be a vector whose coordinates are independent, distributed as x and let Γ be the measurement matrix generated by x . Note that up to the normalization factor of $\|z\|_{L_2}$, which is of the order of a constant when $R^2\delta \lesssim 1$, $\sqrt{N}\Gamma$ is a perturbation of a Rademacher matrix by a sparse matrix with few random spikes that are either R or $-R$.

As noted earlier, if for every $t \in \mathbb{R}^n$,

$$\|\langle X, t \rangle\|_{L_4} \leq C \|\langle X, t \rangle\|_{L_2}, \quad (3.2)$$

then the small-ball condition holds with constants that depend only on C . To show that X satisfies (3.2), denote by \mathbb{E}_η (resp. \mathbb{E}_ε) the expectation with respect to the η -variables (resp. ε -variables), and observe that by a straightforward application of Khintchine’s inequality (see, e.g., p.91 in [27]), for every $t \in \mathbb{R}^n$,

$$\begin{aligned} \mathbb{E}\langle X, t \rangle^4 &\lesssim \mathbb{E}_\eta \mathbb{E}_\varepsilon \left(\sum_{j=1}^n \varepsilon_j (1 + R\eta_j) t_j \right)^4 \lesssim \mathbb{E}_\eta \left(\sum_{j=1}^n (1 + R\eta_j)^2 t_j^2 \right)^2 \\ &= \mathbb{E}_\eta \sum_{k, \ell} (1 + R\eta_k)^2 t_k^2 (1 + R\eta_\ell)^2 t_\ell^2 \lesssim \|t\|_2^4 = \left(\mathbb{E}\langle X, t \rangle^2 \right)^2, \end{aligned}$$

provided that $R^4\delta \lesssim 1$.

Set $(f_i)_{i=1}^N$ to be the canonical basis of \mathbb{R}^N and put $\tilde{\Gamma} = (z_{\ell k}) = \|z\|_{L_2} \sqrt{N}\Gamma$, an $N \times n$ matrix whose entries are independent copies of z . Let

$$v_j = \tilde{\Gamma}e_j = \sum_{\ell=1}^N z_{\ell j}f_\ell,$$

and consider

$$V = \text{absconv}(\{v_j : 2 \leq j \leq n\}),$$

the convex hull of $(\pm v_j)_{j=2}^n$.

We will show that with probability at least $1/2$, $\sqrt{N}B_2^N \subset V$ and $\|v_1\|_2 \leq \sqrt{N}$, in three steps:

Lemma 3.2 *With probability at least $3/4$, for every $1 \leq i \leq N$ there is $y_i \in B_\infty^N$ for which $y_i + Rf_i \in V$.*

In other words, with non-trivial probability, V contains a perturbation of all the vectors $Rf_i, i = 1, \dots, N$, and thus, V ‘almost’ contains RB_1^N .

Proof. Fix a realization of the $N \times n$ Rademacher matrix $(\varepsilon_{\ell j})$ and note that for every $1 \leq i \leq N$ and every $2 \leq j \leq n$

$$v_j = \sum_{\ell=1}^N \varepsilon_{\ell j}f_\ell + \varepsilon_{ij}Rf_i$$

if $\eta_{ij} = 1$ and for every $\ell \neq i$, $\eta_{\ell j} = 0$. Moreover, if this happens, and since V is centrally symmetric (that is, if $v \in V$ then $-v \in V$),

$$\varepsilon_{ij} \left(\sum_{\ell=1}^N \varepsilon_{\ell j}f_\ell \right) + Rf_i = y_i + Rf_i \in V,$$

and $y_i \in B_\infty^N$.

Thus, it remains to estimate the probability that for every $1 \leq i \leq N$ there is some $2 \leq j \leq n$ for which $\eta_{ij} = 1$ and for every $\ell \neq i$, $\eta_{\ell j} = 0$. Clearly, for every $1 \leq i \leq N$,

$$\begin{aligned} & P_\eta \left(\text{there exists } j \in \{2, \dots, n\} : \eta_{ij} = 1, \text{ and } \eta_{\ell j} = 0 \text{ if } \ell \neq i \right) \\ &= 1 - (1 - (1 - \delta)^{N-1}\delta)^{n-1} \geq 1 - \frac{1}{4N} \end{aligned}$$

provided that

$$\frac{\log N}{n} \lesssim \delta \lesssim \frac{\log(en/N)}{N}.$$

Hence, the claim follows by the union bound and integration with respect to the (ε_{ij}) . \blacksquare

Next, it is straightforward to verify that when V contains such a perturbation of RB_1^N (by vectors in B_∞^N), it must also contain a large Euclidean ball, assuming that R is large enough.

Lemma 3.3 *Let $R > N$, and for every $1 \leq i \leq N$, set $y_i \in B_\infty^N$ and put $v_i = Rf_i + y_i$. If V is a convex, centrally symmetric set, and if $v_i \in V$ for every $1 \leq i \leq N$ then $(R/\sqrt{N} - \sqrt{N})B_2^N \subset V$.*

Proof. A separation argument shows that if $\sup_{v \in V} |\langle v, w \rangle| \geq \rho$ for every $w \in S^{N-1}$, then $\rho B_2^N \subset V$ (indeed, otherwise there would be some $x \in \rho B_2^N \setminus V$; but it is impossible to separate x and the convex and centrally symmetric V using any norm-one functional).

To complete the proof, observe that for every $w \in S^{N-1}$,

$$\begin{aligned} \sup_{v \in V} |\langle v, w \rangle| &\geq \max_{1 \leq i \leq N} |\langle Rf_i + y_i, w \rangle| \\ &\geq \max_{1 \leq i \leq N} |\langle Rf_i, w \rangle| - \max_{1 \leq i \leq N} |\langle y_i, w \rangle| \geq R/\sqrt{N} - \sqrt{N}. \end{aligned}$$

\blacksquare

Applying Lemma 3.3, it follows that if $R \geq 2N$ then with probability at least $3/4$, $\sqrt{N}B_2^N \subset V$. Finally, if $\delta \lesssim 1/N$ then

$$P\left(\sum_{\ell=1}^N z_{\ell 1} f_\ell \in \sqrt{N}B_2^N\right) \geq P\left(\left\|\sum_{\ell=1}^N z_{\ell 1} f_\ell\right\|_2 = \sqrt{N}\right) = (1 - \delta)^N \geq 3/4.$$

Hence, with probability at least $1/2$,

$$\sum_{\ell=1}^N z_{\ell 1} f_\ell = \tilde{\Gamma}e_1 \in V = \text{absconv}(\tilde{\Gamma}e_j : j \in \{2, \dots, n\}),$$

and the same assertion holds for the normalized matrix Γ , showing that it does not satisfy ER(1).

Of course, this assertion holds under several conditions on the parameters involved: namely, that $R = \sqrt{p}(1/\delta)^{1/p} \geq 2N$; that $(\log N)/n \lesssim \delta \lesssim \log(en/N)/N$; that $R^4\delta \lesssim 1$; that $p \leq 2\log(1/\delta)$ and that $\delta \lesssim 1/N$.

For instance, one may select $\delta \sim (\log N)/n$ and $p \sim (\log n)/\log N$, in which case all these conditions are met; hence, with probability at least $1/2$, Γ does not satisfy ER(1), proving Theorem C. A similar calculation leads to the proof of Theorem C'. ■

Remark 3.4 *Note that the construction leads to a stronger, non-uniform result, namely, that for every basis vector e_k , with probability at least $1/2$, e_k is not the unique solution of $\min(\|t\|_1 : \Gamma t = \Gamma e_k)$. In particular, uniformity over all supports of size 1 in the definition of ER(1) is not the reason why the moment assumption in Theorem A is required.*

4 Results in the noisy measurements setup

In previous sections, we considered the idealized scenario, in which the data was noiseless. Here, we will study the noisy setup: one observes N couples $(z_i, X_i)_{i=1}^N$, and each z_i is a noisy observation of $\langle X_i, x_0 \rangle$:

$$z_i = \langle X_i, x_0 \rangle + g_i, \quad i = 1, \dots, N. \quad (4.1)$$

The goal is to obtain as much information as possible on the unknown vector x_0 with only the data $(z_i, X_i)_{i=1}^N$ at one's disposal, and for the sake of simplicity, we will assume that the g_i 's are independent Gaussian random variables $\mathcal{N}(0, \sigma^2)$ that are also independent of the X_i 's.

Unlike the noiseless case, there is no hope of reconstructing x_0 from the given data, and instead of exact reconstruction, there are three natural questions that one may consider:

- The *estimation problem*: given some norm $\|\cdot\|$ on \mathbb{R}^n , one would like to construct a procedure \hat{x} for which $\|\hat{x} - x_0\|$ is as small as possible.
- The *prediction problem*: given a new (random, independent) ‘input’ $X \in \mathbb{R}^n$, one has to find a good guess $\langle \hat{x}, X \rangle$ of the most likely associated output z , knowing that (z, X) shares the same distribution with the other couples $(z_1, X_1), \dots, (z_N, X_N)$.
- The *de-noising problem*: given a norm $\|\cdot\|$ on \mathbb{R}^N and a measurement matrix Γ , one has to construct \hat{x} for which $\|\Gamma \hat{x} - \Gamma x_0\|$ is small.

These three problems are central in modern Statistics, and are featured in numerous statistical monographs, particularly in the context of the Gaussian regression model (Equation (4.1)).

Recently, all three problems have been recast in a ‘high-dimensional’ scenario, in which the number of observations N may be much smaller than the ambient dimension n . Unfortunately, such problems are often impossible to solve without additional assumptions, and just as in the noiseless case, the situation improves dramatically if x_0 has some *low-dimensional structure*, for example, if it is s -sparse. The aim is therefore to design a procedure that performs as if the true dimension of the problem is s rather than n , despite the noisy data.

To that end, ℓ_0 penalization methods, sometimes called Model Selection procedures, have been introduced and studied extensively (see, e.g., [29, 4] for results in the context of the model (4.1), as well as in other examples). However, just as in the noise-free problem, the obvious downside of ℓ_0 penalization methods is that they are not feasible computationally. This has led to the introduction of convex relaxations, based on ℓ_1 minimization.

Two well established ℓ_1 -based procedures are the *LASSO* (see, e.g., [43]) defined by

$$\hat{x}_\lambda \in \operatorname{argmin}_{x \in \mathbb{R}^n} \left(\frac{1}{N} \sum_{i=1}^N (z_i - \langle X_i, x \rangle)^2 + \lambda \|x\|_1 \right), \quad (4.2)$$

and the Dantzig selector (see [11]).

Both procedures may be implemented effectively, and their estimation and de-noising properties have been obtained under some assumptions on the measurement matrix (see, e.g. [6, 3, 44] or Chapters 7 and 8 in [24]).

In this section, we shall focus on two such conditions on the measurement matrix. The first, called the Compatibility Condition, was introduced in [44] (see also Definition 2.1 in [45]); the second, the Restricted Eigenvalue Condition, was introduced in [3].

Definition 4.1 *Let Γ be an $N \times n$ matrix. For $L > 0$ and a set $S \subset \{1, \dots, n\}$, the **compatibility constant** associated with L and S is*

$$\phi(L, S) = \sqrt{|S|} \min \left(\|\Gamma \zeta_S - \Gamma \zeta_{S^c}\|_2 : \|\zeta_S\|_1 = 1, \|\zeta_{S^c}\|_1 \leq L \right), \quad (4.3)$$

where ζ_S (resp. ζ_{S^c}) denotes a vector that is supported in S (resp. S^c).

Γ satisfies the **Compatibility Condition for the set S_0** with constants $L > 1$ and c_0 if $\phi(L, S_0) \geq c_0$; it satisfies the **uniform Compatibility Condition (CC) of order s** if $\min_{|S| \leq s} \phi(L, S) \geq c_0$.

A typical result for the LASSO in the Gaussian model (4.1) and when Γ satisfies the Compatibility Condition, is Theorem 6.1 in [6]:

Theorem 4.2 ([6], Theorem 6.1) *Let $x_0 \in \mathbb{R}^n$ be some fixed vector and assume that the data $(z_i, X_i)_{i=1}^N$ have been drawn according to the Gaussian regression model (4.1). Denote by $\Gamma = N^{-1/2} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i$ the measurement matrix. Let $t > 0$. If S_0 is the support of x_0 and $\lambda = 4\sigma \sqrt{(t^2 + \log n)/N}$, then with probability larger than $1 - 2 \exp(-t^2/2)$,*

$$\|\Gamma \hat{x}_\lambda - \Gamma x_0\|_2^2 \leq \frac{64\sigma^2 \|x_0\|_0 (t^2 + \log n)}{N\phi^2(3, S_0)} \quad \text{and} \quad \|\hat{x}_\lambda - x_0\|_1 \leq \frac{64\sigma \|x_0\|_0}{\phi^2(3, S_0)} \sqrt{\frac{t^2 + \log n}{N}}.$$

Even though the Compatibility Condition in S_0 suffices to show that the LASSO is an effective procedure, the fact remains that S_0 is not known. And while a non-uniform approach is still possible (e.g., if Γ is a random matrix, one may try showing that with high probability it satisfies the Compatibility Condition for the fixed, but unknown S_0), the uniform Compatibility Condition is a safer requirement – and the one we shall explore below.

Another uniform condition of a similar flavour is the Restricted Eigenvalue Condition from [3]. To define it, let us introduce the following notation: for $x \in \mathbb{R}^n$ and a set $S_0 \subset \{1, \dots, n\}$ of cardinality $|S_0| \leq s$, let S_1 be the subset of indices of the m largest coordinates of $(|x_i|)_{i=1}^n$ that are outside S_0 . Let $x_{S_{01}}$ be the restriction of x to the set $S_{01} = S_0 \cup S_1$.

Definition 4.3 *Let Γ be an $N \times n$ matrix. Given $c_0 \geq 1$ and an integer $1 \leq s \leq m \leq n$ for which $m + s \leq n$, the **restricted eigenvalue constant** is*

$$\kappa(s, m, c_0) = \min \left(\frac{\|\Gamma x\|_2}{\|x_{S_{01}}\|_2} : S_0 \subset \{1, \dots, n\}, |S_0| \leq s, \|x_{S_0^c}\|_1 \leq c_0 \|x_{S_0}\|_1 \right).$$

*The matrix Γ satisfies the **Restricted Eigenvalue Condition (REC)** of order s with a constant c if $\kappa(s, s, 3) \geq c$.*

Estimation and de-noising results follow from Theorem 6.1 (for the Dantzig selector) and Theorem 6.2 (for the LASSO) in [3], when the measurement matrix Γ , normalized by having the diagonal elements of $\Gamma^\top \Gamma$ equal 1, satisfies the REC of an appropriate order and with a constant that is independent of the dimension. We also refer to Lemma 6.10 in [6] for similar results that do not require normalization.

Because the two lead to bounds on the performance of the LASSO and the Dantzig selector, a question that comes to mind is whether there are matrices that satisfy the CC or the REC. And, as in Compressed Sensing,

the only matrices that are known to satisfy those conditions for the optimal number of measurements (rows) are well-behaved random matrices (see [35, 39, 34, 46] for some examples).

Our aim in this final section is to extend our results to the noisy setup, by identifying almost necessary and sufficient moment assumptions for the CC and the REC. This turns out to be straightforward: on one hand, the proof of Theorem A actually provides a stronger quantitative version of the exact reconstruction property; on the other, the uniform compatibility condition can be viewed as a *quantitative version* of a geometric condition on the polytope ΓB_1^n that characterizes Exact Reconstruction. A similar observation is true for the REC: it can be viewed as a quantitative version of the null space property (see [18, 19] and below) which is also equivalent to the exact reconstruction property.

Definition 4.4 *Let $1 \leq s \leq N$. A centrally symmetric polytope $P \subset \mathbb{R}^N$ is ***s-neighbourly*** if every set of s of its vertices, containing no antipodal pair, is the set of all vertices of some face of P .*

It is well known [16] that Γ satisfies ER(s) if and only if ΓB_1^n has $2n$ vertices and ΓB_1^n is a centrally symmetric s -neighbourly polytope. It turns out that this property is characterized by the uniform CC.

Lemma 4.5 *Let Γ be an $N \times n$ matrix. The following are equivalent:*

1. ΓB_1^n has $2n$ vertices and is s -neighbourly,
2. $\min(\phi(1, S) : S \subset \{1, \dots, n\}, |S| \leq s) > 0$.

In particular, $\min_{|S| \leq s} \phi(L, S)$ for some $L \geq 1$ is a quantitative measure of the s -neighbourly property of ΓB_1^n : if ΓB_1^n is s -neighbourly and has $2n$ vertices then the two sets

$$\{\Gamma \zeta_S : \|\zeta_S\|_1 = 1\} \text{ and } \{\Gamma \zeta_{S^c} : \|\zeta_{S^c}\|_1 \leq 1\} \quad (4.4)$$

are disjoint for every $|S| \leq s$. However, $\min_{|S| \leq s} \phi(1, S)$ measures how far the two sets are from one another, uniformly over all subsets $S \subset \{1, \dots, n\}$ of cardinality at most s .

Proof. Let C_1, \dots, C_n be the n columns of Γ . It follows from Proposition 2.2.13 and Proposition 2.2.16 in [12] that ΓB_1^n has $2n$ vertices and is a centrally symmetric s -neighbourly polytope if and only if for every $S \subset \{1, \dots, n\}$ of cardinality $|S| \leq s$ and every choice of signs $(\varepsilon_i) \in \{-1, 1\}^S$,

$$\text{conv}(\{\varepsilon_i C_i : i \in S\}) \cap \text{absconv}(\{C_j : j \notin S\}) = \emptyset. \quad (4.5)$$

It is straightforward to verify that

$$\bigcup_{(\varepsilon_i) \in \{\pm 1\}^S} \text{conv}(\{\varepsilon_i C_i : i \in S\}) = \{\Gamma \zeta_S : \|\zeta_S\|_1 = 1\}$$

and that

$$\text{absconv}(\{C_j : j \notin S\}) = \{\Gamma \zeta_{S^c} : \|\zeta_{S^c}\|_1 \leq 1\}.$$

As a consequence, (4.5) holds for every $S \subset \{1, \dots, n\}$ of cardinality at most s if and only if $\min(\phi(1, S) : S \subset \{1, \dots, n\}, |S| \leq s) > 0$. ■

An observation of a similar nature is true for the REC: it can be viewed as a quantitative measure of the null space property.

Definition 4.6 *Let Γ be an $N \times n$ matrix. Γ satisfies the **null space property of order s** if it is invertible in the cone*

$$\{x \in \mathbb{R}^n : \text{there exists } S \subset \{1, \dots, n\}, |S| \leq s \text{ and } \|x_{S^c}\|_1 \leq \|x_S\|_1\}. \quad (4.6)$$

In [18, 19], the authors prove that Γ satisfies ER(s) if and only if it has the null space property of order s .

A natural way of quantifying the invertibility of Γ in the cone (4.6) is to consider its smallest singular value, restricted to this cone, which is simply the REC $\kappa(s, n - s, 1)$. Unfortunately, statistical properties of the LASSO and of the Dantzig selector are not known under the assumption that $\kappa(s, n - s, 1)$ is an absolute constant (though if $\kappa(s, s, 3)$ is an absolute constant, LASSO is known to be optimal [3]).

The main result of this section is the following:

Theorem E. *Let $L > 0$, $1 \leq s \leq n$ and $c_0 > 0$. Under the same assumptions as in Theorem A and with the same probability estimate, $\Gamma = N^{-1/2} \sum_{i=1}^N \langle X_i, \cdot \rangle f_i$ satisfies:*

1. *A uniform compatibility condition of order $c_1 s$, namely that*

$$\min_{|S| \leq c_1 s} \phi(L, S) \geq u^2 \beta / 4$$

for $c_1 = u^2 \beta / (16e(1 + L)^2)$.

2. *A restricted eigenvalue condition of order $c_2 s$, with*

$$\kappa(c_2 s, m, c_0) \geq u^2 \beta / 4$$

for any $1 \leq m \leq n$, as long as $(1 + c_0)^2 c_2 \leq u^2 \beta / (16e)$.

On the other hand, if Γ is the matrix considered in Theorem C, then with probability at least $1/2$, $\phi(1, \{e_1\}) = 0$ and $\kappa(1, m, 1) = 0$ for any $1 \leq m \leq n$.

Just like Theorem A and Theorem C, Theorem E shows that the requirement that the coordinates of the measurement vector have $\log n$ moments is almost a necessary and sufficient condition for the uniform Compatibility Condition and the Restricted Eigenvalue Condition to hold. Moreover, it shows the significance of the small-ball condition, even in the noisy setup.

It also follows from Theorem E that if X satisfies the small-ball condition and its coordinates have $\log n$ well-behaved moments as in Theorem A, then ΓB_1^n has $2n$ vertices and is s -neighbourly with high probability for $N \sim s \log(en/s)$. In particular, this improves Theorem 4.3 in [1] by a logarithmic factor for matrices generated by subexponential variables.

Proof of Theorem E: Fix a constant c_1 to be named later and let $S \subset \{1, \dots, n\}$ of cardinality $|S| \leq c_1 s$. Let $\zeta_S \in \mathbb{R}^n$ be a vector supported on S with $\|\zeta_S\|_1 = 1$ and let $\zeta_{S^c} \in \mathbb{R}^n$ be supported on S^c with $\|\zeta_{S^c}\|_1 \leq L$.

Consider $\gamma = (\zeta_S - \zeta_{S^c}) / \|\zeta_S - \zeta_{S^c}\|_2$. Since

$$\|\zeta_S - \zeta_{S^c}\|_2 \geq \|\zeta_S\|_2 \geq \frac{\|\zeta_S\|_1}{\sqrt{|S|}} = \frac{1}{\sqrt{|S|}},$$

it follows that $\gamma \in ((1+L)\sqrt{|S|})B_1^n \cap S^{n-1}$.

Recall that by (2.7), if $r = (1+L)^2 c_1 s \leq su^2\beta/(16e)$, then $\|\Gamma\gamma\|_2 \geq (u^2\beta)/4$. Therefore,

$$\|\Gamma\zeta_S - \Gamma\zeta_{S^c}\|_2 \geq \frac{u^2\beta}{4} \|\zeta_S - \zeta_{S^c}\|_2 \geq \frac{u^2\beta}{4} \|\zeta_S\|_2 \geq \frac{u^2\beta \|\zeta_S\|_1}{4\sqrt{|S|}} = \frac{u^2\beta}{4\sqrt{|S|}},$$

and thus $\min_{|S| \leq c_1 s} \phi(L, S) \geq u^2\beta/4$ for $c_1 = u^2\beta/(16e(1+L)^2)$.

Turning to the REC, fix a constant c_2 to be named later. Consider x in the cone and let $S_0 \subset \{1, \dots, n\}$ of cardinality $|S_0| \leq c_2 s$ for which $\|x_{S_0^c}\|_1 \leq c_0 \|x_{S_0}\|_1$. Let $S_1 \subset \{1, \dots, n\}$ be the set of indices of the m largest coordinates of $(|x_i|)_{i=1}^n$ that are outside S_0 and put $S_{01} = S_0 \cup S_1$.

Observe that $\|x\|_1 \leq (1+c_0) \|x_{S_0}\|_1 \leq (1+c_0)\sqrt{|S_0|} \|x\|_2$; hence $x/\|x\|_2 \in ((1+c_0)\sqrt{|S_0|})B_1^n \cap S^{n-1}$. Applying (2.7) again, if $(1+c_0)^2 c_2 s \leq su^2\beta/(16e)$, then $\|\Gamma x\|_2 \geq ((u^2\beta)/4) \|x\|_2$. Thus,

$$\frac{\|\Gamma x\|_2}{\|x_{S_{01}}\|_2} \geq \frac{\|\Gamma x\|_2}{\|x\|_2} \geq \frac{u^2\beta}{4}$$

and $\kappa(c_2 s, m, c_0) \geq u^2\beta/4$ for any $1 \leq m \leq n$, as long as $(1+c_0)^2 c_2 \leq u^2\beta/(16e)$.

The proof of the second part of Theorem E is an immediate corollary of the construction used in Theorem C. Recall that with probability at least $1/2$, $\Gamma e_1 \in \text{absconv}(\Gamma e_j : j \in \{2, \dots, n\})$. Setting $J = \{e_2, \dots, e_n\}$, there is $\zeta \in B_1^J$ for which $\|\Gamma e_1 - \Gamma \zeta\|_2 = 0$. Therefore, $\phi(1, \{e_1\}) = 0$ and $\kappa(1, m, 1) = 0$ for any $1 \leq m \leq n$, as claimed. \blacksquare

Remark 4.7 *The results obtained in Theorem A and in parts (1) and (2) of Theorem E are also valid for the normalized (columns wise) measurement matrix:*

$$\Gamma_1 = \Gamma \tilde{D}^{-1} \text{ where } \tilde{D} = \text{diag}(\|\Gamma e_1\|_2, \dots, \|\Gamma e_n\|_2).$$

The proof is almost identical to the one used for Γ itself, even though Γ_1 does not have independent rows vectors, due to the normalization. For the sake of brevity, we will not present the straightforward proof of this observation.

Finally, the counterexample constructed in the proof of Theorem C and in which a typical Γ does not satisfy $ER(1)$, does not necessarily generate ΓB_1^n that is not s -neighbourly. Indeed, an inspection of the construction shows that the reason $ER(1)$ fails is that ΓB_1^n has less than $2n - 2$ vertices, rather than that ΓB_1^n is not s -neighbourly. Thus, the question of whether a moment condition is necessary for the random polytope ΓB_1^n to be s -neighbourly with probability at least $1/2$ is still unresolved.

References

- [1] Radosław Adamczak, Alexander E. Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann. Restricted isometry property of matrices with independent columns and neighborly polytopes by random sampling. *Constr. Approx.*, 34(1):61–88, 2011.
- [2] Keith Ball. Cube slicing in \mathbf{R}^n . *Proc. Amer. Math. Soc.*, 97(3):465–473, 1986.
- [3] Peter J. Bickel, Ya’acov Ritov, and Alexandre B. Tsybakov. Simultaneous analysis of LASSO and Dantzig selector. *Ann. Statist.*, 37(4):1705–1732, 2009.
- [4] Lucien Birgé and Pascal Massart. Gaussian model selection. *J. Eur. Math. Soc. (JEMS)*, 3(3):203–268, 2001.
- [5] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013. ISBN 978-0-19-953525-5.
- [6] Peter Bühlmann and Sara A. van de Geer. *Statistics for high-dimensional data*. Springer Series in Statistics. Springer, Heidelberg, 2011. Methods, theory and applications.
- [7] Emmanuel J. Candès. The restricted isometry property and its implications for compressed sensing. *C. R. Math. Acad. Sci. Paris*, 346(9-10):589–592, 2008.

- [8] Emmanuel J. Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52(2):489–509, 2006.
- [9] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1207–1223, 2006.
- [10] Emmanuel J. Candès and Terence Tao. Near-optimal signal recovery from random projections: universal encoding strategies? *IEEE Trans. Inform. Theory*, 52(12):5406–5425, 2006.
- [11] Emmanuel J. Candès and Terence Tao. The Dantzig selector: statistical estimation when p is much larger than n . *Ann. Statist.*, 35(6):2313–2351, 2007.
- [12] Djilil Chafaï, Olivier Guédon, Guillaume Lecué, and Alain Pajor. *Interactions between compressed sensing random matrices and high dimensional geometry*, volume 37 of *Panoramas et Synthèses [Panoramas and Syntheses]*. Société Mathématique de France, Paris, 2012.
- [13] Scott Shaobing Chen, David L. Donoho, and Michael A. Saunders. Atomic decomposition by basis pursuit. *SIAM Rev.*, 43(1):129–159, 2001. Reprinted from *SIAM J. Sci. Comput.* **20** (1998), no. 1, 33–61 (electronic) [MR1639094 (99h:94013)].
- [14] Jon F. Claerbout and Francis Muir. Robust modeling with erratic data. *Geophysics*, 38(5):826–844, 1973.
- [15] Víctor H. de la Peña and Evarist Giné. *Decoupling*. Probability and its Applications (New York). Springer-Verlag, New York, 1999. From dependence to independence, Randomly stopped processes. U -statistics and processes. Martingales and beyond.
- [16] David L. Donoho. Neighborly polytopes and sparse solutions of under-determined linear equations. Technical report, Department of Statistics, Stanford University, 2005.
- [17] David L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.
- [18] David L. Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via l^1 minimization. *Proc. Natl. Acad. Sci. USA*, 100(5):2197–2202 (electronic), 2003.
- [19] David L. Donoho and Xiaoming Huo. Uncertainty principles and ideal atomic decomposition. *IEEE Trans. Inform. Theory*, 47(7):2845–2862, 2001.
- [20] David L. Donoho and Ben F. Logan. Signal recovery and the large sieve. *SIAM J. Appl. Math.*, 52(2):577–591, 1992.
- [21] R. M. Dudley. Central limit theorems for empirical measures. *Ann. Probab.*, 6(6):899–929 (1979), 1978.
- [22] Simon Foucart. Stability and robustness of ℓ_1 -minimizations with Weibull matrices and redundant dictionaries. *Linear Algebra Appl.*, 441:4–21, 2014.
- [23] Simon Foucart and Holger Rauhut. *A mathematical introduction to compressive sensing*. Applied and Numerical Harmonic Analysis. Birkhäuser/Springer, New York, 2013.

- [24] Vladimir Koltchinskii. *Oracle inequalities in empirical risk minimization and sparse recovery problems*, volume 2033 of *Lecture Notes in Mathematics*. Springer, Heidelberg, 2011. Lectures from the 38th Probability Summer School held in Saint-Flour, 2008, École d'Été de Probabilités de Saint-Flour. [Saint-Flour Probability Summer School].
- [25] Vladimir Koltchinskii and Shahar Mendelson. Bounding the smallest singular value of a random matrix without concentration. Technical report, Technion and Georgia Tech, 2013. arXiv:1312.3580.
- [26] Rafał Łatała. Estimation of moments of sums of independent real random variables. *Ann. Probab.*, 25(3):1502–1513, 1997.
- [27] Michel Ledoux and Michel Talagrand. *Probability in Banach spaces*, volume 23 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1991. Isoperimetry and processes.
- [28] Ben Logan. *Properties of High-Pass Signals*. PhD thesis, Columbia University, New York, 1965.
- [29] Pascal Massart. *Concentration inequalities and model selection*, volume 1896 of *Lecture Notes in Mathematics*. Springer, Berlin, 2007. Lectures from the 33rd Summer School on Probability Theory held in Saint-Flour, July 6–23, 2003, With a foreword by Jean Picard.
- [30] Shahar Mendelson. Learning without concentration. *Journal of the ACM*. To appear.
- [31] Shahar Mendelson. A remark on the diameter of random sections of convex bodies. *Geometric Aspects of Functional Analysis (GAFA Seminar Notes)*. Lecture notes in Mathematics 2116, pages 395–404, 2014.
- [32] Shahar Mendelson, Alain Pajor, and Nicole Tomczak-Jaegermann. Uniform uncertainty principle for Bernoulli and subgaussian ensembles. *Constr. Approx.*, 28(3):277–289, 2008.
- [33] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24(2):227–234, 1995.
- [34] Roberto Imbuzeiro Oliveira. The lower tail of random quadratic forms, with applications to ordinary least squares and restricted eigenvalue properties. Technical report, IMPA, 2013. arXiv:1312.2903.
- [35] Garvesh Raskutti, Martin J. Wainwright, and Bin Yu. Restricted eigenvalue properties for correlated Gaussian designs. *J. Mach. Learn. Res.*, 11:2241–2259, 2010.
- [36] B. A. Rogozin. An estimate for the maximum of the convolution of bounded densities. *Teor. Veroyatnost. i Primenen.*, 32(1):53–61, 1987.
- [37] Mark Rudelson and Roman Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Comm. Pure Appl. Math.*, 61(8):1025–1045, 2008.
- [38] Mark Rudelson and Roman Vershynin. Small ball probabilities for linear images of high dimensional distributions. Technical report, University of Michigan, 2013.
- [39] Mark Rudelson and Shuheng Zhou. Reconstruction from anisotropic random measurements. *IEEE Trans. Inform. Theory*, 59(6):3434–3447, 2013.
- [40] Fadil Santosa and William W. Symes. Linear inversion of band-limited reflection seismograms. *SIAM J. Sci. Statist. Comput.*, 7(4):1307–1330, 1986.

- [41] Nikhil Srivastava and Roman Vershynin. Covariance estimation for distributions with $2 + \varepsilon$ moments. *Ann. Probab.*, 41(5):3081–3111, 2013.
- [42] Howard L. Taylor, Stephen C. Banks, and John F. McCoy. Deconvolution with the l_1 -norm. *Geophysics*, 44(1):39–52, 1979.
- [43] Robert Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288, 1996.
- [44] Sara A. van de Geer. The deterministic lasso. In *JSM proceedings. American Statistical Association*, 140, 2007.
- [45] Sara A. van de Geer. Weakly decomposable regularization penalties and structured sparsity. Technical report, ETH Zürich, 2013.
- [46] Sara A. van de Geer and Alan Muro. On higher order isotropy conditions and lower bounds for sparse quadratic forms. Technical report, ETH Zürich, 2014.
- [47] Aad W. van der Vaart and Jon A. Wellner. *Weak convergence and empirical processes*. Springer Series in Statistics. Springer-Verlag, New York, 1996. With applications to statistics.