Spectral Properties of Heavy-Tailed Random Matrices

by

Elizaveta Rebrova

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in the University of Michigan
2018

Doctoral Committee:

Professor Jinho Baik, Co-Chair
Professor Roman Vershynin, Co-Chair, University of California, Irvine
Professor Alexander Barvinok
Professor Raj Rao Nadakuditi
Professor Mark Rudelson

Elizaveta Rebrova

erebrova@umich.edu

ORCID iD: 0000-0002-4041-4238

© Elizaveta Rebrova 2018

Acknowledgments

I have many people to thank for this thesis.

First, I would like to thank my advisor, Prof. Roman Vershynin, for his support and advice throughout all five years of my PhD, for the collaboration, and for the chance to learn from his unique, beautiful and highly motivating way of doing mathematics.

I would like to thank my co-advisor, Prof. Jinho Baik, for being very supportive and helpful as my last year advisor, and for the thesis presentation suggestions. I am very grateful to Prof. Mark Rudelson for organizing lots of invaluable learning opportunities, and for suggesting the majority of the great conferences I attended during my graduate career. I am grateful to Prof. Alexander Barvinok for the suggestions on improvement of this thesis. I would like to thank Kostya Tikhomirov for the inspiring collaboration and friendship. I am grateful to all the people who were hosting me, listened to my talks and discussed mathematics with me. I would like to specially thank Professors Alexander Litvak, Nicole Tomczak-Jaegermann, Olivier Guédon, Pierre Youssef, Afonso S. Bandeira and Raj Rao Nadakuditi.

I am grateful to all the people and organizations that supported me during these five years financially, including University of Michigan Department of Mathematics, Rackham Graduate School, NSF, and Smilka Zdravkovska. I would also like to thank the staff at the Math Department, and especially Tara McQueen, for always being patient and helpful. Thanks to my Ann Arbor friends for the all the student seminars, coffee, food and memories we shared. For lots of inspiring conversation both about mathematics and life, and for making me feel really welcome in Ann Arbor I would like to thank Prof. Sergey Fomin and Irina Arakelian.

I am grateful to the people who have shown me the beauty of mathematics and the taste of tackling hard problems. In particular, to my mom, to my high school math teacher Lev Altshuler, and to the memory of another great teacher Vitaly Arnold, I was lucky to learn from you. I am infinitely thankful to all my family for always staying close despite of living so far, for almost unconditional belief in my abilities, and for supporting me through all ups and downs of the last five years. Finally, I cannot thank enough my husband Dmitry for an amazing gift of love and exploration of the life together.

TABLE OF CONTENTS

LIST OF FIGURES

ABSTRACT

The classical Random Matrix Theory studies asymptotic spectral properties of random matrices when their dimensions grow to infinity. In contrast, the non-asymptotic branch of the theory is focused on explicit high probability estimates that we can obtain for large enough, but fixed size random matrices. This goal naturally brings into play some beautiful methods of high-dimensional probability and geometry, such as the concentration of measure phenomenon.

One of the less understood random matrix models is a heavy-tailed model. This is the case when the matrix entries have distributions with slower tail decay than gaussian, e.g., with a few finite moments only. This work is devoted to the study of the heavy-tailed matrices and addresses two main questions: invertibility and regularization of the operator norm.

First, the invertibility result of Rudelson and Vershynin is generalized from the case when the matrix entries are subgaussian to the case when only two finite moments are required. Then, it is shown that the operator norm of a matrix can be reduced to the optimal order $O(\sqrt{n})$ if and only if the entries have zero mean and finite variance. We also study the constructive ways to perform such regularization. We show that deletion of a few large entries regularizes the operator norm only if all matrix entries have more than two finite moments. In the case with exactly two finite moments, we propose an algorithm that zeroes out a small fraction of the matrix entries to achieve the operator norm of an almost optimal order $O(\sqrt{\ln\ln n \cdot n})$. Finally, if in the latter case the matrix has scaled Bernoulli entries, we get a stronger regularization algorithm that provides a) $O(\sqrt{n})$-operator norm of the resulting matrix and b) simple structure of the "bad" submatrix to be zeroed out.

CHAPTER 1

Introduction

1.1 Non-asymptotic random matrix theory

Random matrix theory is a beautiful area where the methods of probability, functional analysis, linear algebra, and combinatorics can be applied together to study the structure and properties of the matrices taken from some probability distribution. Classical and powerful way to understand the structure of the random matrix is to look at the matrix spectrum: eigenvalues, eigenvectors, or singular values and vectors. The crucial observation that makes results of the theory possible is that the spectrum stabilizes as the size of the matrices (taken from some distribution over all the matrices) grows to infinity, so there are limit laws for the distribution of the spectrum. Among the classical limit laws are, for example, classical Wigner semicircular law for the limiting eigenvalue empirical measure distribution of random symmetric matrices, Tracy–Widom law for the limiting distribution of the largest eigenvalue, Marchenko–Pastur law for the sample covariance matrices. A number of great surveys and books are written on the subject, including [Tao12, AGZ10, Bor17].

But what if instead of the estimates in the limit we would like to get an explicit high probability estimate for a large, but finite size random matrices? This is what non-asymptotic branch of the theory concerned with.

One powerful tool for showing that some properties hold with high probability is the concentration of measure phenomenon. It is a geometrical property of high dimensional measure spaces, which states – in its simplest form – that the majority of the volume of the unit sphere in $\mathbb{R}^n$, when n is large, lies in a thin strip near the equator. In more general versions, the concentration of measure guarantees that some complex expressions involving many independent random variables typically stay very close to a constant (expected value). For detailed discussions of the concentration of measure phenomenon applications to random matrices see, e.g., [Ver12, Tro15]. We also revisit some definitions and facts relevant to this work later in Section 2.3. This idea bridges the study of random matrices with the geometry of high dimensional spaces, which turns out to be important in the non-asymptotic random matrix theory both as a tool and as a source of new interesting questions.

Geometric methods and explicit probability estimates for large finite matrices make this area useful in application dealing with large high dimensional data objects (see, for example, [Ver16, LLV18a]). Among the examples are graph and network analysis, statistics, compressed sensing and more.

1.2 Overview of the results.

In this work we study random matrices with the elements in $\mathbb{R}$ via their real spectrum, i.e. singular values and vectors. Singular values

$$s_{max}(A) = s_1(A) \geq s_2(A) \geq \ldots \geq s_n(A) = s_{min}(A) \geq 0$$

of a real $m \times n$ matrix A are the eigenvalues of the matrix $\sqrt{X^T X}$. The largest and the smallest singular values are especially informative as they determine the basic geometric properties of A as a linear operator, namely, the norm of A and its inverse:

$$s_{max}(A) = \sup_{\|x\|_2=1} \|Ax\|_2 =: \|A\|, \tag{1.1}$$

$$s_{min}(A) = \inf_{\|x\|_2=1} \|Ax\|_2 = 1/\|A^{-1}\|. \tag{1.2}$$

Additionally, the *condition number* of the random matrix A is defined as the ratio of the extreme singular values:

$$\sigma(A) := s_{max}(A)/s_{min}(A).$$

Good (i.e., not too large) condition number means that the matrix A would not distort the metric too much, which can be important for geometric applications. From the other point of view, bounded condition number implies the stability of a system of linear equations $Ax = b$ for any right hand side $b \in \mathbb{R}^m$. Clearly, to bound the condition number, we have to estimate the first singular value from above, and the last one from below.

In the first part of the work presented, we estimate (from below) the smallest singular value of the heavy-tailed random matrices.

1.2.1 Invertibility

The study of the smallest singular value of the random matrix A can answer the questions like "what is the probability that A is singular?" or "if this probability is small (and thus A is likely to be invertible), how large is the typical operator norm of the inverse?"

The invertibility problem was extensively studied from several angles, including the methods of additive combinatorics ([TV10a, TV10b]), mathematical physics ([ESY10, ESY09]), and geometric functional analysis. The latter one is the core of our approach. Let us briefly list the previous work leading to our invertibility theorem for the heavy-tailed matrices (Theorem 1.1), and then state the theorem. More complete overview of the related work is given later in Section 3.1.

For the square matrices with independent standard Gaussian entries, limiting distribution of the smallest singular value was computed in [Ede88] and [Sza91]: it was proved that for any $\varepsilon \in (0,1)$

$$\mathbb{P}\{s_n(A) \leq \varepsilon n^{-1/2}\} \leq C\varepsilon,$$

where $C = \sqrt{2e}$ (see [[Sza91], Theorem 1.2]). This result was later extended for more general distributions and improved in a number of papers, including [Rud08, TV09] (random sign matrices), [RV08] (finite fourth moment entries), [TV10b, TV08] (finite second moment entries and non-zero mean), [Tik17] (entries with bounded density). In the paper [RV08] the authors also got much more precise probability estimate for a narrower class of subgaussian matrices. A random variable ξ is called *subgaussian*, if it has at least gaussian tail decay, i.e., there exists a number $K > 0$ such that

$$\mathbb{P}\{|\xi| > t\} \leq 2\exp(-t^2/K^2), \qquad \text{for all } t > 0.$$

In particular, it implies that all the moments of ξ are finite, $\mathbb{E}\xi^p \leq C\sqrt{p}$, where $C = C(K)$ is a constant (see also [Ver12]). In [RV08], Rudelson and Vershynin have shown that an $n \times n$ random matrix A with centered subgaussian random entries satisfies for every $\varepsilon > 0$

$$\mathbb{P}\big\{s_n(A) \leq \varepsilon n^{-1/2}\big\} \leq L\varepsilon + u^n, \tag{1.3}$$

for some constants $L > 0$ and $u \in (0,1)$. Note that this implies linear probability decay with ε (like in the gaussian case) until ε reaches an exponentially small probability level. The latter is necessary since singularity probability for some (e.g. discrete) matrices is non-zero. So, right hand side in (1.3) cannot always decay to zero with $\varepsilon \to 0$ as it does for gaussian matrices. For example, in the case of i.i.d. symmetric Rademacher entries (every entry is ± 1 with probability 0.5) the probability that two rows are equal up to sign is already greater than 0.5^n. Clearly, this lower bounds singularity probability of such matrices.

Our work with K. Tikhomirov [RT15] shows that exactly the same probability estimate holds for much more general class of *heavy-tailed matrices*, such that

$$A \text{ is } n \times n; \text{ the entries of } A \text{ are i.i.d., with } \mathbb{E}A_{ij} = 0, \ \ \mathbb{E}A_{ij}^2 = 1. \tag{1.4}$$

Namely, we proved the following

Theorem 1.1. *Let $n \geq n_0$ and let A be a heavy-tailed matrix satisfying* (1.4). *Then for any $\varepsilon > 0$*

$$\mathbb{P}\big\{s_n(A) \leq \varepsilon n^{-1/2}\big\} \leq L\varepsilon + u^n, \tag{1.5}$$

where $L > 1$ and $u \in (0, 1)$ are constants, depending only on distribution of A_{ij}.

Another result of the same paper – geometric Theorem 1.2, providing a bound for covering numbers for random ellipsoids in high dimensional spaces – naturally appeared as a part of the proof of Theorem 1.1.

The proof of Theorem 1.1, as well as the prior result of Rudelson and Vershynin, is based on a version of *ε-net argument*. One of the basic applications of the net argument is as follows. Let us approximate the unit sphere $S^{n-1} := \{\|x\|_2 = 1\}$ by a finite set of points (net) in order to reduce complexity of the infimum in (1.2), and then use concentration of measure techniques to conclude the result for every $x \in S^{n-1}$ (see, for example, [Ver12]). This standard form of the argument is classic and elegant, and also too weak for the square matrices.

In [RV08] a special net refinement (depending on the arithmetic structure of a random vector) was introduced. In [RT15] we use yet another net refinement on the image of the unit ball under the action of A. This led us to construct a $O(\sqrt{n})$-net for a random ellipsoid $\{Ax : \|x\|_2 = 1\}$. We proved the following geometric theorem:

Theorem 1.2. *For any $\delta < 1/4$ and $n > n_0(\delta)$ there exists a subset $\mathcal{N} \subset B_2^n$ of cardinality at most $(6/\delta)^{13n\delta}$, such the for any $n \times n$ random matrix A satisfying* (1.4)

$$A(B_2^n) \subset \bigcup_{y \in A(\mathcal{N})} \big(y + \frac{C\sqrt{n}}{\delta} B_2^n\big)$$

with probability at least $1 - 4\exp(-\delta n/8)$. Here $C > 0$ is a universal constant.

This result can also be interpreted geometrically in terms of covering numbers. The *covering number* $N(S, K)$ for two subsets S and K of a vector space is defined as the smallest number of parallel translates of K sufficient to cover S. By Theorem 1.2, we have an exponential bound on the covering number $N(A(B_2^n), \frac{C\sqrt{n}}{\delta}B_2^n) \leq (6/\delta)^{13n\delta}$ with high probability. Moreover, the covering is non-random and *"universal"* for all random ellipsoids of the type $\{Ax : \|x\|_2 = 1\}$, independent from the distribution of the entries of A (as long as A satisfies (1.4)).

1.2.2 Regularization of the norm

In the second part of the work we study the maximal singular value of a random matrix, which is, as noted in (1.1), also its operator norm. Our goal is to impose a good upper estimate on it in the heavy-tailed case.

If we consider a matrix A with independent standard Gaussian entries, then by the classical Bai-Yin law (see, for example, [Tao12])

$$s_1(A)/\sqrt{n} \to 2 \quad \text{almost surely,}$$

as the dimension $n \to \infty$. Moreover, the $2\sqrt{n}$ asymptotics holds for more general classes of matrices. By [YBK88], if the entries of A have zero mean and bounded fourth moment, then

$$\|A\| = (2 + o(1))\sqrt{n}$$

with high probability. In the non-asymptotic regime, an application of Bernstein's inequality (see, for example, in [Ver12]) gives

$$\mathbb{P}\{s_1(A) \leq t\sqrt{n}\} \geq 1 - e^{-c_0 t^2 n} \quad \text{for } t \geq C_0$$

for the matrices with i.i.d. subgaussian entries. Here, $c_0, C_0 > 0$ are absolute constants. The non-asymptotic extensions to more general distributions are also available, see [Seg00, Lat05, BVH16, vH17a].

Note that the order $\sqrt{n}$ is the best we can generally hope for. Indeed, if the entries of A have variance C, then the typical magnitude of the Euclidean norm of a row of A is $\sim \sqrt{n}$, and the operator norm of A cannot be smaller than that. So, it is natural to assume $O(\sqrt{n})$ as the "ideal order" of the operator norm of an $n \times n$ i.i.d. random matrix.

We do not have ideal $O(\sqrt{n})$-order in the heavy-tailed regime (1.4). It is suggested by the fact that weak fourth moment is necessary for the convergence in probability of $\|A\|/\sqrt{n}$ when n grows to infinity (see [Sil89]). Moreover, an explicit family of examples, constructed in [LS14], shows heavy-tailed matrices A that have $\|A\| \sim O(n^\alpha)$ for any $\alpha \leq 1$ with substantial probability.

A natural question is: what are the obstructions in the structure of A that cause too large norm? Can we regularize the matrix restoring the "ideal norm"? Clearly, interesting regularization would be the one that does not change A too much.

So, the first question to answer is when we can enforce the "ideal" $\|A\|$ with high probability by modifying just a small fraction of the entries? Let us assume nothing about the distribution of the i.i.d. entries of A. Under what moment conditions the regularization of the matrix norm is a *local problem*?

We answered these questions in our work with R. Vershynin [RV18]. We have shown that local regularization (on a small submatrix of the matrix) is possible if and only if the entries of A have zero mean and unit variance:

Theorem 1.3 (Local problem). *Consider an $n \times n$ random matrix A satisfying* (1.4)*, and let $\varepsilon \in (0, 1/6]$. Then, with probability at least $1 - 7\exp(-\varepsilon n/12)$, there exists an $\varepsilon n \times \varepsilon n$ submatrix of A such that replacing all of its entries with zero leads to a well-bounded matrix $\tilde{A}$:*

$$\|\tilde{A}\| \leq \frac{C \ln \varepsilon^{-1}}{\sqrt{\varepsilon}} \cdot \sqrt{n},$$

where C is a sufficiently large absolute constant.

Theorem 1.4 (Global problem). *Consider an $n \times n$ random matrix A_n whose entries are i.i.d. copies of a random variable that has either nonzero mean or infinite second moment, and let $\varepsilon \in (0, 1)$. Then*

$$\min \frac{\|\tilde{A}_n\|}{\sqrt{n}} \to \infty \quad \textit{as} \quad n \to \infty$$

almost surely. Here the minimum is with respect to the matrices $\tilde{A}_n$ obtained by any modification of any $\varepsilon n \times \varepsilon n$ submatrix of A_n.

Our proof utilizes the cut norm and Grothendieck-Pietsch factorization for matrices ([Pie78, AN04]), and it combines the methods developed in [RT15] and [LLV17]. See Section 4.1.1 for the proof overview and further discussion.

1.2.3 Constructive regularization and random graphs

Returning to the question about the obstructions: if the local regularization is possible (that is, the matrix satisfies (1.4)), what exactly causes the norm of a centered random matrix A to be too large? Is there a simple description of the small fraction of elements to be deleted, or at least an algorithm that allows to determine them?

This question is not answered in full in this work, but we present several partial results obtained: for the matrices with more than two finite moments, for the adjacency matrices of random graphs, and for the general case.

More than two finite moments

First natural guess would be that the only troublemakers are a few large entries of A, and so we can obtain a result like Theorem 1.3 simply by zeroing them out. This intuition turns out to be misleading. Only in the case when A_{ij} have *more* than two finite moments the truncation idea works and it is not hard to derive the following result from known bounds on random matrices such as [vH17a, Seg00, AT16]:

Theorem 1.5 ($2+\varepsilon$ moments)**.** *Let $\varepsilon \in (0,1]$ and $n > n_0(\varepsilon)$. Consider any $n \times n$ random matrix A with i.i.d. mean zero entries which satisfy $\mathbb{E}|A_{ij}|^{2+\varepsilon} \leq 1$. With probability at least $1 - 2\exp(-n^{\varepsilon/5})$, zeroing out at most $n^{1-\varepsilon/9}$ largest entries of A leads to the matrix $\tilde{A}$ such that*

$$\|\tilde{A}\| \leq 8\sqrt{n}.$$

The heavy-tailed model (1.4) is qualitatively harder: to find a small submatix, deletion of which regularizes the norm, one have to account for the mutual positions of large elements of A.

Sparse random graphs

Another case when it is clear how to regularize the operator norm is when the matrix is Bernoulli. It is closely related to the work of Le, Levina and Vershynin [LLV17] (see also the survey [LLV18a]) on inhomogeneous Erdős-Rényi random graphs. *Inhomogeneous Erdős-Rényi random graph $G(n, p_{ij})$* is a graph with n vertices, such that the edge between i-th and j-th vertices is present with probability p_{ij} (independently of all other edges). Often (for example, in community detection problems), the question of interest is to estimate some features of the probability matrix (p_{ij}) from random graphs drawn from $G(n, p_{ij})$. Concentration of the adjacency matrix A around its expectation matrix $\mathbb{E}A$, when it holds, guarantees that such features can be recovered.

However, this concentration holds only in the case when the graph is *dense* (i.e. maximal expected degree of the vertices $d := \max_{ij} np_{ij} > \log n$). For the *sparse* graphs, especially when expected degree is constant, the regularization question appears. Can we modify the graph on a small subset of its edges such that $\|A - \mathbb{E}A\|$ become well-bounded, and so concentration will be restored? And which exactly modification can help us?

An answer was given in the work of Feige and Ofek [FO05]: with high probability, it is enough to delete all the edges adjacent to the "heavy" vertices (n/d vertices with the largest degrees). Le, Levina and Vershynin presented in [LLV17] a more general way to regularize these edges, as well as a new approach to the proof. Using similar techniques, we can check that it is actually enough to modify a much smaller subgraph of the graph and also described the structure of this "bad" subgraph:

Theorem 1.6. *Let $G = G(n, (p_{ij}))$ be an inhomogeneous Erdős-Rényi graph and let d denotes its average expected degree $d = \max_{ij} np_{ij}$, $d \geq 5$ and let $r \geq 1$. Then for any n large enough with probability at least $1 - 6 \cdot (10ne^{-d})^{-r}$ we can enumerate the vertices of G such that the adjacency matrix A has the following properties:*

- *if $\tilde{A}$ is obtained from A by zeroing out the entries of top left $s \times s$ submatrix A_0, then*

$$\|\tilde{A} - \mathbb{E}A\| \leq Cr^{3/2}\sqrt{d},$$

- *A_0 is a very small square part of A, as its size $s \lesssim 10ne^{-d} \ll n/d$,*

- *A_0 has no more than $40r$ ones in every column above its main diagonal.*

The last condition means that we can direct the edges inside the "bad" (too dense) subgraph such that every vertex will have a finite number of the outcoming edges (see also discussion in the Remark 5.6).

General case

Finally, for general matrices with exactly two finite moments, weaker *constructive* versions of Theorem 1.3 are possible. Based on simple individual correction of the entries, a bound with an additional factor $\log n$ in the norm and weaker probability guarantees can be derived from known general bounds on random matrices, such as the matrix Bernstein's inequality ([Tro15]). One would apply the matrix Bernstein's inequality for the entries truncated at level $\sqrt{n}$ to get that $\|\tilde{A}\| \leq \varepsilon^{-1/2}\sqrt{n} \cdot \ln n$ (this is shown in Lemma 5.13).

More sophisticated regularization procedure, based on finding small fraction of "heavy" rows and columns to delete, can regularize the norm to the order $\sqrt{n \cdot \log \log n}$ (see Theorem 5.14 for the precise statement; the proof is based on the works [FO05, BVH16]).

1.3 Outline of the Dissertation

In Chapter 2 we gather useful definitions and notations, as well as several concentration and discretization lemmas, that are going to be used throughout the text.

In Chapter 3 we prove invertibility Theorem 1.1 and covering Theorem 1.2 as its intermediate step.

Chapters 4 and 5 are devoted to the study of s_{max}, or, equivalently, operator norm of the matrix. In Chapter 4 we prove that zero first and finite second moment is necessary and sufficient condition for a square i.i.d. random matrix A to have a regularize version: an another matrix $\bar{A}$ that differs from A in a small square sub-matrix only and has the operator norm $\|\bar{A}\| \sim \sqrt{n}$ (Theorems 1.3 and 1.4).

Discussion of a possibility of constructive regularization starts in the Section 4.4. Then, in Chapter 5 we discuss various results related to the constructive regularization: when A has more than two finite moments (Section 5.1); when A is a Bernoulli matrix (Section 5.2); and for the general case with extra $\sqrt{\log \log n}$ term (Section 5.3). Also, in the Section 5.2 we discuss the interpretation of Bernoulli matrices as adjacency matrices of the Erdős-Réniy random graphs and the structure of a small "bad" sub-graph (to be deleted in regularization).

Finally, in Chapter 6 we mention several open questions that are closely connected to the research presented, and seem interesting for the future exploration.

CHAPTER 2

Preliminaries

2.1 Notations

Throughout the work, positive absolute constant are denoted C, C_1, c, c_1, etc. Their values may be different from line to line. However, sometimes, to avoid confusion, we add a numerical subscript to the name of a constant defined within a statement (e.g. constant $C_{1.57}$ is the one that appears in Lemma 1.57).

We often write $a \lesssim b$ to indicate that $a \leq Cb$ for some absolute constant C.

Given a finite set S, by $|S|$ we denote its cardinality. By $e_1, e_2, \ldots, e_n$ we denote the canonical basis in $\mathbb{R}^n$. The standard inner product in $\mathbb{R}^n$ shall be denoted by $\langle \cdot, \cdot \rangle$. Given $p \in [1, \infty]$, $\|\cdot\|_p$ is the standard ℓ_p-norm. For ℓ_2, will will simply write $\|\cdot\|$. Given an $m \times n$ matrix M and $p, q \in [1, \infty]$, by $\|M\|_{p\to q}$ we denote the operator norm of M considered as the mapping from $(\mathbb{R}^n, \|\cdot\|_p)$ to $(\mathbb{R}^m, \|\cdot\|_q)$.

The discrete interval $\{1, 2, \ldots, n\}$ is denoted by $[n]$. If $\mathcal{R}$ is some subset of indices, $\mathcal{R} \subset [n] \times [n]$, let us denote by $A_{\mathcal{R}}$ the matrix obtained from A by replacing by zero the entries A_{ij} with the indices $(i, j) \in \mathcal{R}^c$:

$$A_{\mathcal{R}} := (\bar{A}_{ij})_{i,j=1}^n, \text{ where } \bar{A}_{ij} = A_{ij}\mathbb{1}_{\{(i,j)\in\mathcal{R}\}}.$$

We will often consider subsets of columns of the matrix, so when $\mathcal{R} = J \times [n]$ we use a simplified notation: for $J \subset [n]$

$$A_J := A_{[n]\times J}.$$

Finally, main probability model to be considered is the following:

$$A \text{ is } n \times n\text{; the entries of } A \text{ are i.i.d., with } \mathbb{E}A_{ij} = 0, \ \ \mathbb{E}A_{ij}^2 = 1. \tag{*}$$

We will sometimes refer to it as probability model (*), and explicitly describe all the changes in the model we consider.

2.2 Nets and covering numbers

In this section we give necessary definitions to be used later and briefly preview the role of nets and covering numbers in this thesis. A detailed discussion of the subject can be found, e.g., in [[Ver16], Section 4.2].

Definition 2.1 (ε-net)**.** Let (T, d) be a metric space. For any subset $K \subset T$ and $\varepsilon > 0$, a subset $\mathcal{N} \subset K$ is called an ε-net of K if for every $x \in K$ there exists $x_0 \in \mathcal{N}$ such that $d(x, x_0) \leq \varepsilon$. Equivalently, $\mathcal{N}$ is an ε-net of K if and only if K can be covered by balls with centers in $\mathcal{N}$ and radii ε.

Usually, nets in $\mathbb{R}^n$ are defined with respect to Euclidean metric, i.e., for every x, y in $\mathbb{R}^n$ define $d(x, y) = \|x - y\|$. However, other metrics can be useful in special situations. Moreover, Definition 2.1 generalizes without changes to pseudometrics. We will also consider nets with respect to the pseudometric $d_A(x, y) := \|A(x - y)\|$, where x and y are points in $\mathbb{R}^n$ and A is a linear operator (see Section 3.2.1).

Definition 2.2 (Covering numbers)**.** The smallest possible cardinality of an ε-net of K is called the covering number of K and is denoted $\mathbf{N}(K, \varepsilon)$. Equivalently, covering number is the smallest number of closed balls with centers in K and radii ε whose union covers K.

Almost equivalent definition of the *exterior covering number* $\mathbf{N}^{ext}(K, \varepsilon)$ allows the centers of ε balls to lie outside of K. It is not hard to check that

$$\mathbf{N}^{ext}(K, \varepsilon) \leq \mathbf{N}(K, \varepsilon) \leq \mathbf{N}^{ext}(K, \varepsilon/2).$$

Finally, we can define the covering number for two general subsets $S, K \subset T$ as the smallest number of parallel translates of K sufficient to cover S, denoted as $\mathbf{N}(S, K)$.

It is a well-known fact (e.g. [[Ver16], Corollary 4.2.13]) that a unit ball in $\mathbb{R}^n$ have ε-covering number at most

$$\mathbf{N}(B_2^n, \varepsilon) \leq \left(\frac{3}{\varepsilon}\right)^n. \tag{2.1}$$

In particular, this fact is used to bound the $s_1(A) = \sup_{x \in S^{n-1}} \|Ax\|$ for the random matrix A with independent subgaussian entries. General idea is as follows: for every fixed $x \in S^{n-1}$ the norm of $\|Ax\|$ is tightly bounded using subgaussian concentration, and then union bound is taken over the ε-net on S^{n-1} with the cardinality at most $(3/\varepsilon)^n$.

However, in many interesting cases, including smallest singular value of subgaussian matrices, the estimate (2.1) is too weak, so that union bound does not preserve good estimates, holding for individual $x \in S^{n-1}$. This brings us to the need to consider *net refinements* – more complicated ε-nets of lower cardinality, specified to the subsets of S^{n-1}, possessing additional nice properties.

It is probably not surprising that for the heavy-tailed matrices we need even more delicate net estimates: the tight subgaussian concentration is no longer available, so even individual estimates on $\|Ax\|$ with some fixed x are weaker. Net refinement idea is crucial for the proofs of Theorems 1.1 and 1.2 (see also the discussion in the Section 3.2.1).

2.3 Concentration

A standard way to get some desired estimate on a random variable X *with high probability* is to get this estimate for $\mathbb{E}X$ first, and then argue that X *concentrates* around its expectation. In this case X usually stays close to $\mathbb{E}X$, and therefore satisfies a close estimate.

This work relies a lot on the good concentration properties of sums of subgaussian (and sub-exponential) random variables, that is, such that grow not faster than standard normal (respectively, exponential) random variables.

Definition 2.3 (subgaussian random variable)**.** A random variable Y is called *subgaussian* if its moments satisfy

$$\mathbb{E}\exp(Y^2/M_2^2) \leq e,$$

for some number $M_2 > 0$. The minimal number M_2 is called the subgaussian moment of X, denoted as $\|Y\|_{\psi_2}$.

Equivalently, ξ is *subgaussian* if there exists a number $K > 0$ such that

$$\mathbb{P}\{|\xi| > t\} \leq 2\exp\bigl(-t^2/K^2\bigr), \quad t > 0. \tag{2.2}$$

It is easy to check that the smallest value of K satisfying (2.2) is equivalent to *the subgaussian norm* of ξ (see, for example, [Ver12, Lemma 5.5]).

Definition 2.4 (sub-exponential random variable)**.** Analogously, a random variable is called *sub-exponential* if

$$\mathbb{E}\exp(Y/M_1) \leq e,$$

for some number $M_1 > 0$. The minimal number M_1 is called the sub-exponential moment of Y, denoted as $\|Y\|_{\psi_1}$.

The class of subgaussian random variables contains standard normal, Bernoulli, and generally all bounded random variables. The class of sub-exponential random variables is exactly the class of squares of subgaussians. See [Ver12] for more information and statements of standard concentration inequalities.

2.3.1 Some classical concentration inequality

The following inequalities are going to be used in the latter sections. The first one is a standard Khintchine–type inequality (see, for example, [Hoe63]).

Lemma 2.5 (Khintchine's inequality)**.** *Let $r_1, r_2, \dots, r_n$ be independent Rademacher random variables. Then for any vector $y \in S^{n-1}$ the random variable $\sum_{i=1}^n y_i r_i$ is $C_{2.5}$-subgaussian, where $C_{2.5} > 0$ is a universal constant.*

The sum of squares of subgaussian variables has good concentration properties; the bound below follows from a standard "Laplace transform" argument (see, for example, [Ver12, Corollary 5.17]):

Lemma 2.6 (Subexponential tails inequality)**.** *For any $T > 0$ there is $L_{2.6} > 0$ depending on T with the following property: Let $\xi_1, \xi_2, \dots, \xi_n$ be independent centered 1-subgaussian random variables. Then*

$$\mathbb{P}\Big\{\sum_{i=1}^n \xi_i^2 > L_{2.6} n\Big\} \le \exp(-Tn).$$

For the bounded distributions we will use Bernstein's inequality (see, for example, [Ver16, Theorem 2.8.4]):

Lemma 2.7 (Bernstein's inequality for bounded distributions)**.** *Let $\xi_1, \dots, \xi_n$ be independent, mean zero random variables, such that $|\xi_k| \le K$ almost surely for all i. Then, for every $t \ge 0$, we have*

$$\mathbb{P}\left\{\left|\sum_{i=1}^N \xi_i\right| \ge t\right\} \le 2\exp\left(-\frac{t^2/2}{\sigma^2 + Kt/3}\right).$$

Here $\sigma^2 = \sum_{i=1}^N \mathbb{E}\xi_i^2$ is the variance of the sum.

There are concentration inequalities that are applicable directly to sums of independent random matrices, rather than individual random variables, like the following (see, e.g. [Ver16, Theorem 5.4.1])

Lemma 2.8 (Matrix Bernstein inequality)**.** *Let $X_1, \dots, X_N$ be independent, mean zero, $n \times n$ symmetric random matrices, such that $\|X_i\| \le K$ almost surely for all i. Then, for every $t \ge 0$, we have*

$$\mathbb{P}\left\{ \left\| \sum_{i=1}^N X_i \right\| \ge t \right\} \le 2n \exp\left(-\frac{t^2/2}{\sigma^2 + Kt/3} \right).$$

Here $\sigma^2 = \left\| \sum_{i=1}^N \mathbb{E} X_i^2 \right\|$ is the norm of the matrix variance of the sum.

The matrix Bernstein inequality can be extended to non-symmetric (or even non-square) random matrix X, by applying it to a symmetric block matrix, having zero matrices in diagonal blocks and X and X^T in the off-diagonal. Resulting statement is the following:

Corollary 2.9 (Matrix Bernstein inequality for non-symmeric case)**.** *Let $X_1, \dots, X_N$ be independent, mean zero, $m \times n$ random matrices, such that $\|X_i\| \le K$ almost surely for all i. Then, for every $t \ge 0$, we have*

$$\mathbb{P}\left\{ \left\| \sum_{i=1}^N X_i \right\| \ge t \right\} \le 2(n+m) \exp\left(-\frac{t^2/2}{\sigma^2 + Kt/3} \right).$$

Here

$$\sigma^2 = \max\left(\left\| \sum_{i=1}^N \mathbb{E} X_i^T X_i \right\|, \left\| \sum_{i=1}^N \mathbb{E} X_i X_i^T \right\| \right).$$

Finally, for a special case of Bernoulli random variables, one can apply the following classical inequality that provides a quite sharp result, sensitive to the expectations of the individual variables (see, e.g. [Ver16, Theorem 2.3.1])

Lemma 2.10 (Chernoff's inequality)**.** *Let ξ_i be independent Bernoulli random variables with expectations $\mathbb{E}\xi_i = p_i$. Consider their sum $S_N = \sum_{i=1}^N \xi_i$ and denote its mean by $\mu = \mathbb{E} S_N$. Then, for any $t > \mu$, we have*

$$\mathbb{P}\{S_N \ge t\} \le e^{-\mu} \left(\frac{e\mu}{t} \right)^t.$$

2.3.2 Concentration inequality for random permutations

We will also need a concentration inequality for random permutations:

Lemma 2.11 (Concentration for random permutations)**.** *Consider two arbitrary vectors $a = (a_1, \dots, a_n) \in \mathbb{R}^n$ and $x \in \{-1, 1\}^n$. Let $\pi : [n] \to [n]$ denote a random permutation* *chosen uniformly from the symmetric group S_n. Then the random sum*

$$S := \sum_{i=1}^{n} a_i x_{\pi(i)}$$

is subgaussian, and

$$\|S - \mathbb{E}S\|_{\psi_2} \leq C_{2.11} \|a\|_2.$$

The same inequality holds for the sum $S' = \sum_{i=1}^{n} a_{\pi(i)} x_i$ as well, since it has the same distribution as S.

For the proof of this lemma we will need the following definition (essentially taken from [MS86]). Let S be a finite set and d be a pseudometric on S. We say that (S, d) is *of length at most ℓ* (for some $\ell > 0$) if there is $n \in \mathbb{N}$, positive numbers $b_1, b_2, \dots, b_n$ with $\|(b_1, b_2, \dots, b_n)\| \leq \ell$ and a sequence $(S_k)_{k=0}^{n}$ of partitions of S such that

1. $S_0 = \{S\}$;
2. $S_n = \{\{s\}\}_{s \in S}$;
3. S_k is a refinement of S_{k-1} for all $k = 1, 2, \dots, n$;
4. For each $k \in \{1, 2, \dots, n\}$ and any $Q, Q' \in S_k$ such that $Q \cup Q'$ is a subset of an element of S_{k-1}, there is a one-to-one mapping $\phi : Q \to Q'$ such that $d(s, \phi(s)) \leq b_k$ for all $s \in Q$.

In particular, the above conditions on S_k imply that all elements of S_k have the same cardinality.

Theorem 2.12 (see [MS86, Theorem 7.8])**.** *Let (S, d) be a finite pseudometric space of length at most ℓ and let μ be the normalized counting measure on S. Then for any function $f : S \to \mathbb{R}$ satisfying $|f(s) - f(s')| \leq d(s, s')$ ($s, s' \in S$) and all $t > 0$ we have*

$$\mu\Big\{\Big|f - \int f \, d\mu\Big| \geq t\Big\} \leq 2 \exp\Big(-\frac{t^2}{4\ell^2}\Big).$$

Remark 2.13. In [MS86], the above theorem is formulated for metric spaces. It is easy to see that passing to pseudometrics does not change the picture.

Denote by S_n the set of permutations of $[n] := \{1, 2, \dots, n\}$.

Lemma 2.14. *Let $a = (a_1, a_2, \dots, a_n)$ be a non-zero vector and $x = (x_1, x_2, \dots, x_n)$ be a vertex of the cube $[-1, 1]^n$. Further, let μ be the normalized counting measure on S_n.*

Define a function $f : S_n \to \mathbb{R}$ *as*

$$f(p) := \sum_{j=1}^{n} x_{p(j)} a_j, \quad p \in S_n.$$

Then

$$\mu\Big\{\Big|f - \int f\,d\mu\Big| \geq t\Big\} \leq 2\exp\Big(-\frac{t^2}{64\|a\|_2^2}\Big), \quad t > 0.$$

Proof. Without loss of generality, we can assume that $|a_j| \geq |a_{j+1}|$ ($j = 1, 2, \ldots, n-1$). Define a pseudometric d on S_n: for any $p, q \in S_n$ let

$$d(p, q) := |f(p) - f(q)|.$$

Further, we define a sequence of partitions $(S_{n,k})_{k=0}^{n}$ of S_n: let $S_{n,0} := \{S_n\}$ and for each $k = 1, 2, \ldots, n$, let $S_{n,k}$ consist of all subsets of S_n of the form

$$\{p \in S_n : \ p(1) = i_1, p(2) = i_2, \ldots, p(k) = i_k\}$$

for all $\{i_1, i_2, \ldots, i_k\} \subset [n]$.

Now, let $k \in \{1, 2, \ldots, n\}$ and let $Q, Q' \in S_{n,k}$ be such that $Q \cup Q'$ is a subset of an element of $S_{n,k-1}$. Note that there are numbers $i_1, i_2, \ldots, i_k$, i'_k such that $p(j) = i_j$ for all $j < k$ and $p \in Q \cup Q'$; $p(k) = i_k$ for all $p \in Q$ and $p(k) = i'_k$ for all $p \in Q'$. Define a one-to-one mapping $\phi : Q \to Q'$ by

$$\phi(p)(j) := p(j) \text{ for } j \neq k, p^{-1}(i'_k); \ \ \phi(p)(k) := i'_k; \ \ \phi(p)(p^{-1}(i'_k)) := i_k.$$

For any $p \in Q$, we have

$$d(p, \phi(p)) \leq 2|a_k| + 2|a_{p^{-1}(i'_k)}| \leq 4|a_k|,$$

with the last inequality due to the fact that $p^{-1}(i'_k) \geq k$. Thus, the space (S_n, d) is of length at most $4\|a\|_2$. Applying Theorem 2.12, we get the result. □

Proof of Lemma 2.11. Note that Lemma 2.14 directly implies Lemma 2.11 with $C_{2.11} = 8$ when a is not identically zero, and if $a = 0$, the statement of Lemma 2.11 is trivial. □

2.4 Discretization

Two-point distributions, such as Bernoulli or scaled Bernoulli, have an obvious advantage of being a relatively "simple" example of random variables: having just two possible values, they are relatively easy for direct computations. Also, stronger concentration inequalities (like Chernoff's inequality, Lemma 2.10) are available in the case when random variables in question are Bernoulli.

On the other hand, two-point distributions are already sophisticated enough to model heavy-tailed random variables. Consider ξ such that

$$\xi = \begin{cases} 1, & \text{with probability } p; \\ 0, & \text{otherwise.} \end{cases}$$

For small p, say, $p \sim n^{-1}$, we have $\mathbb{E}\xi = p \sim 0$ and $\mathbb{E}\xi^2 \sim p$. Centralizing and scaling ξ to have second moment 1, we arrive to a random variable $\psi \sim p^{-1/2}\xi$ such that for any $\varepsilon > 0$

$$\lim_{n\to\infty} \mathbb{E}\psi^{2+\varepsilon} = \lim_{n\to\infty} n^{\varepsilon/2} = \infty,$$

that is, a random variable with exactly two moments bounded (by a constant independent from n). Clearly, varying the order of success probability p we can obtain "simple" random variables with any certain number of finite moments.

In light of discussion above, sometimes it is convenient to represent a general a. s. non-negative random variable via several scaled Bernoulli random variables. One straightforward way to do it is to observe that

$$X \le \sum_{k=0}^{K} q_k \xi_k, \quad \xi_k := \mathbb{1}_{\{X \in [q_{k-1}, q_k)\}},$$

for any increasing sequence $\{q_k\}$, such that $X \in [q_0, q_K)$ almost surely (or $K = \infty$ and $\{q_k\} \to_k \infty$). Similarly we can approximate from below

$$X \ge \sum_{k=0}^{K} q_{k-1} \mathbb{1}_{\{X \in (q_{k-1}, q_k)\}}.$$

This approximation is used, for example, to prove global obstructions for the operator norm regularization (Section 4.3). One of the important disadvantages of this approximation is that random variables ξ_k are not independent.

2.4.1 Approximation by independent Bernoulli

The following lemma allows us to approximate a general continuous random variable by a sum of independent, scaled Bernoulli random variables.

Lemma 2.15 (Discretization)**.** *Consider a non-negative, continuous random variable X. There exists a non-negative random variable X' satisfying the following.*

1. $\mathbb{E}X' \le 4\mathbb{E}X$.

2. *X' stochastically dominates X, i.e.*

$$\mathbb{P}\{X' \ge t\} \ge \mathbb{P}\{X \ge t\} \quad \textit{for all } t \ge 0.$$

3. *X' is a sum of scaled, independent Bernoulli random variables:*

$$X' = \sum_{k=0}^{\infty} q_k \xi_k \tag{2.3}$$

where q_k are non-negative numbers and ξ_k are independent $\mathrm{Ber}(2^{-k})$ *random variables.*

Proof. Set the values q_k to be the quantiles of the distribution of X:

$$q_k := \min\left\{t \ge 0 : \mathbb{P}\{X \ge t\} = 2^{-k-1}\right\}, \quad k = 0, 1, 2, \ldots$$

(These values are well defined since the cumulative distribution function of X is continuous by assumption.) By definition, (q_k) is an increasing sequence. Define X' by (2.3).

To check part 1, note that by definition,

$$\mathbb{E}X' = \sum_{k=0}^{\infty} q_k \mathbb{E}\xi_k = \sum_{k=0}^{\infty} q_k 2^{-k}. \tag{2.4}$$

To lower bound $\mathbb{E}X$, let us decompose X according to the values it can take. This gives

$$X \ge \sum_{k=0}^{\infty} X \mathbb{1}_{\{X \in [q_k, q_{k+1})\}} \ge \sum_{k=0}^{\infty} q_k \mathbb{1}_{\{X \in [q_k, q_{k+1})\}}$$

almost surely. Taking expectation of both sides, we obtain

$$\mathbb{E}X \ge \sum_{k=0}^{\infty} q_k \mathbb{P}\{X \in [q_k, q_{k+1})\}.$$

Now, using the definition of q_k, we have

$$\mathbb{P}\left\{X \in [q_k, q_{k+1})\right\} = \mathbb{P}\left\{X \geq q_k\right\} - \mathbb{P}\left\{X \geq q_{k+1}\right\} = 2^{-k-1} - 2^{-k-2} = 2^{-k-2}.$$

This yields

$$\mathbb{E}X \geq \sum_{k=0}^{\infty} q_k 2^{-k-2}. \tag{2.5}$$

Comparing (2.4) with (2.5), we conclude that $\mathbb{E}X' \leq 4\mathbb{E}X$, which proves part 1 of the lemma.

Let us prove part 2. If $t \in [q_k, q_k+1)$ for some $k = 0, 1, 2, \ldots$, then using the definitions of X' and q_k we obtain

$$\begin{aligned}\mathbb{P}\left\{X' \geq t\right\} \geq \mathbb{P}\left\{X' \geq q_{k+1}\right\} \geq \mathbb{P}\left\{\xi_{k+1} = 1\right\} = 2^{-k-1} \\ = \mathbb{P}\left\{X \geq q_k\right\} \geq \mathbb{P}\left\{X \geq t\right\},\end{aligned}$$

as required.

It remains to check the domination inequality when t is outside the range $[q_0, q_\infty)$ where $q_\infty := \lim_{k\to\infty} q_k \in \mathbb{R}_+ \cup \{\infty\}$. If $t < q_0$, we have

$$\mathbb{P}\left\{X' \geq t\right\} \geq \mathbb{P}\left\{X' \geq q_0\right\} \geq \mathbb{P}\left\{\xi_0 = 1\right\} = 1,$$

and the inequality in part 2 follows. If $t \geq q_\infty$ then, using the continuity of the cumulative distribution of X, we obtain

$$\mathbb{P}\left\{X \geq t\right\} \leq \mathbb{P}\left\{X \geq q_\infty\right\} = \lim_{k\to\infty} \mathbb{P}\left\{X \geq q_k\right\} = \lim_{k\to\infty} 2^{-k-1} = 0,$$

and the inequality in part 2 follows again. The proof is complete. $\square$

Remark 2.16 (Bounded random variables). Suppose $X \leq M$ almost surely. Then, in the second part of the conclusion of Lemma 2.15, X can be represented as a *finite* sum

$$X' := \sum_{k=0}^{\kappa} q_k \xi_k$$

where q_k are non-negative numbers, $q_k \in [0, M]$, and ξ_k are independent $\mathrm{Ber}(p_k)$ random variables. Here $p_k = 2^{-k} \geq 1/M$ for $k < \kappa$ and $p_\kappa = 1/M$.

Remark 2.17 (Coupling). Stochastic dominance of X' over X in Lemma 2.15 implies that one can realize the random variables X and X' on the same probability space so that

$$X' \geq X \quad \text{almost surely.}$$

(See, for example, [Wol99, Section 4.3], or [Tho00, Chapter 1, Theorem 3.1]).

Moreover, in the same way we can construct a majorizing collection for any collection of independent random variables. In particular, we can do it for all entries of a random matrix A at once.

Lemma 2.15 is used in both invertibility and local obstructions for the norm regularization theorems (Sections 3.2 and 4.2). Fast and controlled decay of the expectations of ξ_k makes it especially convenient for computations.

The discretization X' constructed in Lemma 2.15 is not compatible with original random variable X by size. Namely, we cannot expect that $X \geq c \cdot X'$ for some chosen constant c (although it is true for the expectations: $\mathbb{E}X \geq c \cdot \mathbb{E}X'$). However, it is possible to construct a two-side approximation by a sum of independent scaled Bernoulli. This will be shown in the next section.

2.4.2 Two-side approximation by independent Bernoulli

Approximating "by size" we are bound to have a Bernoulli member of each order, namely, we now take $q_k = 2^k$. Expectations of ξ_k are defined such that there exists a non-zero scaled Bernoulli term of size greater than or equal to t with probability $\mathbb{P}\{X \geq t\}$, for any $t \geq 1$. As a result, we have the following lemma:

Lemma 2.18. *Let $X \geq 0$ a.s. random variable, such that $\mathbb{E}X = 1$, X is not a.s. 1. Then there exists a random variable*

$$\xi = \sum_{i=0}^{\infty} 2^i \xi_i,$$

where ξ_i are independent $0-1$ Bernoulli random variables, such that

$$\xi/2 \leq X < \xi \textit{ a. s. when } X \geq 1.$$

Proof. Without loss of generality we can assume that X has continuous distribution.

Define:

$$\xi_0 = 1 \text{ a.s.}$$

$$\xi_i = 1 \text{ with probability } \frac{\tilde{p}_{i-1} - \tilde{p}_i}{1 - \tilde{p}_i} \text{ and } 0 \text{ otherwise,}$$

where $\tilde{p}_i := \mathbb{P}\{X \geq 2^i\}$. Observe that

$$\mathbb{P}\{X \geq 1\} < 1,$$

hence, all $\tilde{p}_i < 1$ and all ξ_i are well-defined.

We are going to show the stochastic domination, i.e. that

$$\mathbb{P}\{\xi/2 \geq t\} \leq \mathbb{P}\{X \geq t\} < \mathbb{P}\{\xi \geq t\} \text{ for any } t \geq 1.$$

Indeed, consider $t \geq 1$. Let i be such that $t \in (2^{i-1}, 2^i]$. Then

$$\mathbb{P}\{X \geq t\} > \mathbb{P}\{X \geq 2^{i-1}\} = \tilde{p}_{i-1}.$$

On the other hand,

$$\mathbb{P}\{\xi \geq t\} = \mathbb{P}\{\xi \geq 2^i\} =$$

$$\mathbb{P}\{\text{ at least one of the } \xi_j, j \geq i \text{ is non-zero }\} =$$

$$1 - \mathbb{P}\{\text{ all } \xi_j = 0,\ j \geq i\} = 1 - \prod_{j=i}^{\infty} \mathbb{P}\{\xi_j = 0\} =$$

$$1 - \prod_{j=i}^{\infty} \left[1 - \frac{\tilde{p}_{j-1} - \tilde{p}_j}{1 - \tilde{p}_j}\right] = 1 - \prod_{j=i}^{\infty} \left[\frac{1 - \tilde{p}_{j-1}}{1 - \tilde{p}_j}\right] = 1 - (1 - \tilde{p}_{i-1}) = \tilde{p}_{i-1}.$$

Moreover,

$$\mathbb{P}\{\xi/2 \geq t\} = \mathbb{P}\{\xi \geq 2t\} = \mathbb{P}\{\xi \geq 2^{i+1}\} = \tilde{p}_i \leq \mathbb{P}\{X \geq t\}.$$

This completes the proof of Lemma 2.18. □

2.5 Operator norm via ℓ_1 norm of rows and columns

The following simple result, known as the Schur bound [Sch11, p. 6, §2], states that the operator norm of any matrix is dominated by the ℓ_1 norms of rows and columns. For completeness, we state and prove the Schur bound here; the proof is almost identical to the original one.

Lemma 2.19. *For any $m \times k$ matrix A, we have*

$$\|A\| \leq \left(\max_i \|A_i\|_1 \cdot \max_j \|A^j\|_1\right)^{1/2}$$

where A_i and A^j denote the rows and columns of A.

Proof. Recall that the operator norm can be computed as a maximum of the quadratic form:

$$\|A\| = \sup_{\|x\|_2=\|y\|_2=1} |x^\mathsf{T} Ay|.$$

Fix unit vectors x and y and express

$$\begin{aligned}
|x^\mathsf{T} Ay| &= \Big| \sum_{i,j} x_i A_{ij} y_j \Big| \\
&\le \sum_{i,j} \Big(|x_i| \sqrt{|A_{ij}|} \Big) \Big(\sqrt{|A_{ij}|} |y_j| \Big) \quad \text{(by the triangle inequality)} \\
&\le \Big(\sum_{i,j} x_i^2 |A_{ij}| \Big)^{1/2} \Big(\sum_{i,j} |A_{ij}| y_j^2 \Big)^{1/2} \quad \text{(by the Cauchy-Schwarz inequality)} \\
&= \Big(\sum_i x_i^2 \, \|A_i\|_1 \Big)^{1/2} \Big(\sum_j \|A_j\|_1 \, y_j^2 \Big)^{1/2} \\
&\le \max_i \|A_i\|_1^{1/2} \cdot \max_j \|A^j\|_1^{1/2} \quad (\text{since } \|x\|_2 = \|y\|_2 = 1).
\end{aligned}$$

Taking the maximum over all unit vectors x and y, we complete the proof. □

This lemma will be a part of the proof of Theorem 1.3 for moderately large entries of the matrix (in the Section 4.2.3).

Let us briefly summarize what else we saw in this chapter. Nets and covering numbers are main tools in the proof of Theorem 1.1 (recall its intermediate result, Theorem 1.2, which proves an estimate for the covering number for random ellipsoids).

Classical concentration inequalities, gathered in Section 2.3.1, will be of crucial use throughout the whole text of the thesis. Concentration for random permutations (discussed in Section 2.3.2) will be used in both proofs of Theorem 1.1 and Theorem 1.3 for symmetrization purposes, see Proposition 3.15 and Lemma 4.10 respectively.

The discretization by independent Bernoulli (discussed in Section 2.4.1) will be also used for both Theorem 1.1 and Theorem 1.3. It will help us accurately quantify the number of "bad" directions of a random operator defined by a heavy-tailed matrix. See also the discussion in the beginning of Section 3.2.2.

In the following chapter we are going to proceed with the proof of invertibility Theorem 1.1.

CHAPTER 3

Smallest singular value: invertibility

3.1 Motivation and main results

The problem of determining the distribution of the smallest singular value $s_n(A)$ has been given much attention in literature.

Convergence of (appropriately normalized) smallest singular values for a sequence of random rectangular matrices with i.i.d. entries and growing dimensions was established by Bai and Yin [BY08] (see also [Tik15], where the result is proved under optimal moment assumptions). For non-asymptotic results in this direction, we refer the reader to papers [LPRTJ05, RV09] for the case of i.i.d. entries (see also [Tik16] where no moment conditions are assumed); [ALPTJ11, ALPTJ10] for log-concave distributions of rows and [SV13, MP14, KM15, Yas14, GLPTJ17] for more general isotropic distributions. We refer to surveys [RV10, Ver12] (see also [Rud13]) for more information.

For random square matrices with independent standard Gaussian entries, the limiting distribution of the smallest singular value was computed by Edelman [Ede88]; universality of this result was established in [TV10a]. Further, for matrices with i.i.d. entries it was shown in [TV08] and [TV10b] that, given any $K > 0$ there are $R, L > 0$ depending only on K and the law of a_{11} such that $\mathbb{P}\{s_n(A+B) \le n^{-L}\} \le Rn^{-K}$ for any non-random matrix B satisfying $\|B\| \le n^K$ (we note that analogous results were recently obtained for more general models of randomness allowing some dependence between the entries of A; see, in particular, [NO14] and [GNT15]). In the case $B = \mathbf{0}$ which we study, those papers do not provide optimal estimates for $s_n(A)$. A much more precise statement was proved in [RV08] under the additional assumption that the entries of A are subgaussian; namely, Rudelson and Vershynin showed that $s_n(A)$ satisfies a small ball probability estimate

$$\mathbb{P}\{s_n(A) \le \varepsilon n^{-1/2}\} \le L\varepsilon + u^n, \;\; \varepsilon > 0,$$

where $L > 0$ and $u \in (0,1)$ depend only on the subgaussian moment of a_{ij}'s.

The subgaussian condition on the entries is crucial in the proof of this result, as it heavily relies on the assumption that $\|A\| \lesssim \sqrt{n}$ with probability very close to one. In the same paper [RV08], it was shown that $s_n(A) \gtrsim n^{-1/2}$ with large probability provided that the fourth moment of A_{ij} is bounded, however, the statement is much weaker than in the subgaussian case.

Our main goal was to relax as much as possible the moment assumptions on the entries of A while keeping the small ball probability estimate as strong as in the subgaussian theorem of [RV08]. The main result of this chapter is the following

Theorem 3.1. *For any $\widetilde{v} \in (0,1]$ and $\widetilde{u} \in (0,1)$ there are numbers $L > 0$, $u \in (0,1)$ and $n_0 \in \mathbb{N}$ depending only on $\widetilde{v}$ and $\widetilde{u}$ with the following property. Let $n \geq n_0$ and let $A = (a_{ij})$ be an $n \times n$ random matrix with i.i.d. entries such that $\mathbb{E}a_{ij} = 0,\ \ \mathbb{E}a_{ij}^2 = 1$. Let $\sup\limits_{\lambda\in\mathbb{R}} \mathbb{P}\{|a_{11} - \lambda| \leq \widetilde{v}\} \leq \widetilde{u}$. Then for any $\varepsilon > 0$ we have*

$$\mathbb{P}\big\{s_n(A) \leq \varepsilon n^{-1/2}\big\} \leq L\varepsilon + u^n.$$

Note that any random variable α with $\mathbb{E}\alpha = 0$ and $\mathbb{E}\alpha^2 = 1$ obviously satisfies $\sup\limits_{\lambda\in\mathbb{R}} \mathbb{P}\{|\alpha - \lambda| \leq \widetilde{v}\} \leq \widetilde{u}$ for some $\widetilde{v} > 0$ and $\widetilde{u} \in (0,1)$ determined by the law of α. Thus, the above statement does not require any additional assumptions on the matrix apart from (1.4); by introducing the quantities $\widetilde{v}$ and $\widetilde{u}$ we make the dependence of L and u on the law of a_{11} more explicit.

The idea of the proof of Theorem 3.1 can be described as follows. Denote by A' the transpose of the first $n-1$ columns of A. A principal component of the proof of [RV08] is an analysis of the arithmetic structure of null vectors of A', which is described with the help of the notion of the least common denominator (LCD). To show that null vectors of A' typically have an exponentially large LCD, the authors of [RV08] consider subsets S of the unit sphere corresponding to vectors with small LCD, and show that $\inf\limits_{x\in S} \|A'x\| > 0$ with a large probability. For this, they use the standard ε-net argument, when the infimum is estimated by taking a Euclidean ε-net $\mathcal{N}$ on S and applying relation $\inf\limits_{x\in S} \|A'x\| \geq \inf\limits_{y\in\mathcal{N}} \|A'y\| - \varepsilon\|A'\|$ together with the estimate $\|A'\| \leq C\sqrt{n}$ which holds with probability very close to one under the subgaussian moment assumptions on the entries. In our setting, the principal difficulty consists in the fact that the condition (1.4) does not guarantee a good upper bound for the operator norm $\|A'\|$.

To deal with this fundamental issue, we "refine" the nets constructed in [RV08] by the following estimate on a covering number of a random ellipsoid $A(B_2^n)$:

Theorem 3.2. *Let $\delta \in (0, 1/4]$ and $n \geq \frac{1}{4\delta}$. Then there is a (non-random) collection $\mathcal{C}$ of* *parallelepipeds in* $\mathbb{R}^n$ *with* $|\mathcal{C}| \leq \exp(13n\delta \ln \frac{2e}{\delta})$ *having the following property: For any random matrix* A *satisfying* (1.4)*, with probability at least* $1 - 4\exp(-\delta n/8)$ *we have*

$$\forall\, x \in B_2^n \quad \exists P \in \mathcal{C} \text{ such that } x \in P \text{ and } A(P) \subset Ax + \frac{C\sqrt{n}}{\delta} B_2^n.$$

Here, $C > 0$ *is an absolute constant.*

Indeed, it can be shown that Theorem 3.2 implies that, given an ε-net $\mathcal{N}$ on S, it is possible to construct a subset $\widetilde{\mathcal{N}} \subset S$ of cardinality at most $\exp\big(13\delta n \ln \frac{2e}{\delta}\big)|\mathcal{N}|$ which is an $L'\varepsilon\sqrt{n}$-net on S (for some $L' = L'(\delta)$) with respect to the pseudometric $d(x,y) = \|A'(x-y)\|$ with probability at least $1 - 4\exp(-\delta n/8)$. Then, $\inf\limits_{x\in S} \|A'x\| \geq \inf\limits_{y\in\widetilde{\mathcal{N}}} \|A'y\| - L'\varepsilon\sqrt{n}$, so the argument does not depend any more on the value of $\|A'\|$.

The rest of Chapter 3 is organized as follows: Section 3.2 contains a proof of the geometric Theorem 3.2, starting with the discussion of its proof strategy and some interpretations of the result. Then, in Section 3.3.1, we collect some results from [RV08] that are crucial to complete the proof of the main result Theorem 3.1 in Section 3.3.2.

3.2 Coverings of random ellipsoids

3.2.1 Discussion and proof overview

Let us briefly discuss geometric part of the argument: Theorem 3.2 and its more elegant

Corollary 3.3. *For any* $\delta \in (0, 1/4]$ *and* $n \geq \frac{1}{4\delta}$ *there exists a non-random subset* $\mathcal{N} \subset B_2^n$ *of cardinality at most* $\exp(13n\delta \ln \frac{2e}{\delta})$ *such that for any* $n \times n$ *matrix* A *satisfying* (1.4)*, we have*

$$\mathbb{P}\Big\{A(B_2^n) \subset \bigcup_{y\in A(\mathcal{N})} \big(y + \frac{C'\sqrt{n}}{\delta} B_2^n\big)\Big\} \geq 1 - 4\exp(-\delta n/8)$$

for some absolute constant $C' > 0$.

A crucial feature of the results is that the set $\mathcal{C}$ in the theorem is non-random. Moreover, $\mathcal{C}$ (as well as the set $\mathcal{N}$ from Corollary 3.3) provides a "universal" covering which is independent of the distribution of the entries of A. Compared to Corollary 3.3, the statement of Theorem 3.2 is more flexible as it enables us to choose the "anchor" points within the parallelepipeds when constructing corresponding ε-net (this matter is covered in detail at the beginning of Section 3.3.2).

Note that if the entries of A have a bounded fourth moment then the operator norm $\|A\|$ satisfies $\|A\| \le L\sqrt{n}$ with probability close to one (see [YBK88] and [Lat05] for precise statements), whence $\mathbb{P}\{A(B_2^n) \subset L\sqrt{n}B_2^n\} \approx 1$. If, moreover, the entries of A are subgaussian then for some $L > 0$ depending only on the subgaussian moment we have $\mathbb{P}\{A(B_2^n) \subset L\sqrt{n}B_2^n\} \ge 1 - \exp(-n)$. On the other hand, for heavy-tailed entries the operator norm of A may have a higher order of magnitude compared to $\sqrt{n}$ with probability close to one, so the trivial argument given above is not applicable. In our case with only two finite moments, we take advantage of the fact that although large $\|A\|$ implies the existence of "bad" directions $x \in S^{n-1}$, such that $\|Ax\|_2 \gg \sqrt{n}$, with high probability there are only few "bad" directions.

A little more precise description of the main idea of the Theorem 3.2 proof is as follows. The collection $\mathcal{C}$ of parallelepipeds is constructed using a special subset $\mathcal{D}$ of diagonal operators with diagonal elements in the interval $(0, 1]$. Namely, we define $\mathcal{D}$ as the set of all diagonal operators with diagonal entries in $\{1\} \cup \{2^{-2^k}\}_{k=0}^{\infty}$ and with determinants bounded from below by $\exp(-\delta n)$. Then, for every operator D from $\mathcal{D}$, we take a covering of the ball B_2^n by appropriate translates of parallelepiped $D(L''n^{-1/2}B_\infty^n)$ (for some $L'' = L''(\delta)$), and let $\mathcal{C}$ be the union of such coverings over $\mathcal{D}$. It turns out that Theorem 3.2 follows almost immediately from the following relation:

$$\begin{aligned}\mathbb{P}\Big\{&\exists \text{ diagonal matrix } D \text{ with diagonal entries in } \{1\} \cup \{2^{-2^k}\}_{k=0}^{\infty} \text{ such that} \\ &\det D \ge \exp(-\delta n) \text{ and } \|AD\|_{\infty \to 2} \le \frac{Cn}{\sqrt{\delta}}\Big\} \ge 1 - 4\exp(-\delta n/8).\end{aligned} \tag{3.1}$$

In Section 3.2.3, we show that (3.1) holds true under condition (1.4); see Theorem 3.11. Geometrically, this property means that it is possible to construct a random parallelepiped $P \subset [-1, 1]^n$ with sides parallel to the standard coordinate axes, such that $\mathrm{Vol}(P) \ge \exp(-\delta n)$ and A maps P inside the Euclidean ball $\frac{Cn}{\sqrt{\delta}}B_2^n$ with probability at least $1 - 4\exp(-\delta n/8)$. Note that parallelepiped P will be "narrow" along directions $w \in S^{n-1}$ for which $\|Aw\|$ is large.

To put Theorem 3.2 and Corollary 3.3 into more general context, note that they answer, in particular, the following geometrically natural question: *how many translates of the Euclidean ball $\sqrt{n}B_2^n$ (or its constant multiple) are needed to cover the random ellipsoid $A(B_2^n)$?*

Recall that for two subsets S and K of a vector space the *covering number* $\mathbf{N}(S, K)$ is defined as the smallest number of parallel translates of K sufficient to cover S. By Theorem 3.2, $\mathbf{N}(A(B_2^n), \frac{C\sqrt{n}}{\delta}B_2^n) \le \exp\big(13\delta n \ln \frac{2e}{\delta}\big)$ with probability at least $1 - 4\exp(-\delta n/8)$.

Another similar interpretation of the result obtained is in terms of the *entropy number of the operator* A. If X and Y are two Banach spaces, then k-th entropy number of a linear operator $A : X \to Y$ is defined as follows (see, e.g. [Pis99]):

$$e_k := \inf\{\varepsilon > 0 : \mathbf{N}(A(B_X), \varepsilon B_Y) \leq 2^{k-1}\},$$

where B_X and B_Y are the unit balls in X and Y respectively, and k is an integer. By Theorem 3.2, with probability at least $1 - 4\exp(-\delta n/8)$ we have that $e_k(A) \lesssim \sqrt{n}/\delta$ for $k = \lceil 13\sqrt{2}n\delta \ln \frac{2e}{\delta} \rceil$ and any random operator $A : (\mathbb{R}^n, \|.\|) \to (\mathbb{R}^n, \|.\|)$, as long as its matrix in the standard Euclidean basis satisfies probability model (*).

Finally, one more interpretation of these results, that will be of use for us, is related to the net refinement (see Theorem 3.30 in Section 3.3.2). Recall that an *ε-net* $\mathcal{N}$ on a metric space X is a subset of X such that any point of X is within distance at most ε from a point of $\mathcal{N}$ (see also Section 2.2 in Preliminaries). It is easy to see that with probability at least $1 - 4\exp(-\delta n/8)$ the set $\mathcal{N}$ from Corollary 3.3 is a $\frac{C\sqrt{n}}{\delta}$–net on B_2^n with respect to the pseudometric $d(x, y) := \|A(x - y)\|$ $(x, y \in B_2^n)$.

3.2.2 Fitting a random vector into an ℓ_p^n-ball

By $\mathcal{D}_n$ we denote the set of all $n \times n$ diagonal matrices with diagonal elements belonging to the interval $(0, 1]$ (we will sometimes refer to such matrices as positive diagonal contractions). Further, denote by $\mathcal{D}_n^2$ the set of all $n \times n$ positive diagonal contractions whose diagonal entries belong to the set $\{1\} \cup \{2^{-2^k}\}_{k=0}^{\infty}$. The set $\mathcal{D}_n^2$ can be regarded as a discretization of $\mathcal{D}_n$.

In this section, we consider the following problem: Let X be a random vector in $\mathbb{R}^n$ with i.i.d. coordinates. We want to find a random diagonal operator D taking values in $\mathcal{D}_n$ such that $D(X)$ is contained in an appropriate (fixed) multiple of the ℓ_p^n-ball *everywhere* on the probability space and at the same time the determinant of D is typically "not too small". The statement to be proved is

Proposition 3.4. *For any $\alpha \in (0, 1)$ there is a number $L = L(\alpha) > 0$ with the following property: Let $\delta \in (0, 1]$, $p \in [1, \infty)$ and let $X = (x_1, x_2, \ldots, x_n)$ be a random vector on $(\Omega, \Sigma, \mathbb{P})$ with i.i.d. coordinates such that $\mathbb{E}|x_i|^p < \infty$. Then there is a random positive diagonal contraction D taking values in $\mathcal{D}_n$ such that*

$$\|DX\|_p^p \leq \frac{L}{\delta}\mathbb{E}\|X\|_p^p \text{ everywhere on the probability space, and } \mathbb{E}(\det D)^{p\alpha - p} \leq \exp(\delta).$$

Remark 3.5. Proposition 3.4 is a foundation block of the proof. In Section 3.2.3, we will amplify this result (the case $p = 2$) by proving its "matrix version" (Theorem 3.11). The case $p \neq 2$ in this section is considered for completeness. It is not needed later for the proofs of Theorems 3.1 and 3.2.

Remark 3.6. Note that a trivial definition of the diagonal operator $D = (d_{ij})$ by setting

$$d_{jj}{}^p := \min\Big(1, \frac{L}{\delta}\frac{\mathbb{E}\|X\|_p^p}{\|X\|_p^p}\Big), \quad j = 1, 2, \ldots, n,$$

gives an unsatisfactory distribution of the determinant. For example, if the entries of X are $\{0, 1\}$-valued with probability of taking value 1 equal to $1/n$, then $\mathbb{E}\|X\|_p^p = 1$, and for any $m \leq n$ we have

$$\mathbb{P}\{\|X\|_p^p = m\} = \binom{n}{m} n^{-m}\Big(1 - \frac{1}{n}\Big)^{n-m} \geq \frac{1}{4m^m}.$$

Thus, the above definition of D would give $\mathbb{P}\{\det D \leq 2^{-n}\} \geq \frac{1}{4}\lceil 2^p L/\delta\rceil^{-\lceil 2^p L/\delta\rceil}$.

Our construction of the required operator is more elaborate. Let us first describe the idea informally. Assume that $p = 1$ and that X is our random vector with non-negative i.i.d. coordinates with unit expectations. We consider a sequence of non-negative numbers (*levels*) such that each coordinate exceeds k-th level with probability 2^{-k}. The main observation is that X "does not fit" into the ℓ_1^n-ball $\frac{Ln}{\delta}B_1^n$ only if for some k there are much more than $2^{-k}n$ coordinates of X exceeding the level. We define the required operator D so that its restriction to the "bad" coordinates is an appropriate dilation, while on all other coordinates it acts isometrically. If there exist several "bad" levels the operator D will be defined as a product of several diagonal operators. Moreover, it will be more convenient to "replace" the vector X by a sum of independent vectors of two-valued variables (scaled Bernoulli random variables, see also the discussion in the beginning of Section 2.4), such that the sum is a majorant for X on the entire probability space.

The following coupling Lemma 3.7 is a quick corollary of Lemma 2.15. It constructs the discretization by independent scaled Bernoulli random variables for all the elements of a random vector (corresponding to a row of the random matrix) at once.

Lemma 3.7 (Coupling)**.** *Let $Y = (Y_1, Y_2, \ldots, Y_n)$ be a random vector on a probability space $(\Omega, \Sigma, \mathbb{P})$ with i.i.d. non-negative coordinates with everywhere continuous cdf and $\mathbb{E}Y_i = 1$. Then there is a vector $Z = (Z_1, Z_2, \ldots, Z_n)$ on $(\Omega, \Sigma, \mathbb{P})$, and probability space $(\widetilde{\Omega}, \widetilde{\Sigma}, \widetilde{\mathbb{P}})$ and random vectors $\widetilde{Y} = (\widetilde{Y}_1, \widetilde{Y}_2, \ldots, \widetilde{Y}_n)$ and $\widetilde{Z} = (\widetilde{Z}_1, \widetilde{Z}_2, \ldots, \widetilde{Z}_n)$ on $(\widetilde{\Omega}, \widetilde{\Sigma}, \widetilde{\mathbb{P}})$ such that* *a)*

$$Z = \sum_{k=0}^{\infty} \tau_{k+1} \xi^k,$$

where $\xi^k = (\xi_1^k, \xi_2^k, \dots, \xi_n^k)$ *and* ξ_i^k *(*$i \le n$*,* $k = 0, 1, \dots$*) are jointly independent* 0-1 *random variables with* $\mathbb{P}\{\xi_i^k = 1\} = 2^{-k}$ *and* $(\tau_k)_{k=1}^{\infty}$ *is an increasing non-negative sequence satisfying* $\sum_{k=1}^{\infty} \tau_k 2^{-k} < \infty$*.*

b) $\widetilde{Y}$ *and* $\widetilde{Z}$ *are equidistributed with* Y *and* Z*, respectively;*

c) $\widetilde{Z}_i \ge \widetilde{Y}_i$ *for all* $i \le n$ *everywhere on* $(\widetilde{\Omega}, \widetilde{\Sigma}, \widetilde{\mathbb{P}})$*.*

Proof. Fix for a moment $i \le n$ and consider the distributions of Y_i and Z_i. Taking

$$\tau_{k+1} := q_k(Y_i) := \inf\bigl\{t \ge 0 : \mathbb{P}\{Y_i \ge t\} = 2^{-k}\bigr\}, \quad k \ge 0,$$

by Lemma 2.15 and the condition $\mathbb{E}Y_i = 1$, the first condition is satisfied and also Z_i stochastically dominates Y_i.

Then, by the coupling Remark 2.17, there is a probability space $(\widetilde{\Omega}_i, \widetilde{\Sigma}_i, \widetilde{\mathbb{P}}_i)$ and variables $\widetilde{Y}_i$ and $\widetilde{Z}_i$ on $(\widetilde{\Omega}_i, \widetilde{\Sigma}_i, \widetilde{\mathbb{P}}_i)$ equidistributed with Y_i and Z_i, respectively, such that $Z_i \ge Y_i$ everywhere on $\widetilde{\Omega}_i$.

Finally, by taking $(\widetilde{\Omega}, \widetilde{\Sigma}, \widetilde{\mathbb{P}})$ to be the product space $\prod_i \Omega_i$ and naturally extending the variables $\widetilde{Y}_i, \widetilde{Z}_i$ to $(\widetilde{\Omega}, \widetilde{\Sigma}, \widetilde{\mathbb{P}})$, we obtain the random vectors $\widetilde{Y}$, $\widetilde{Z}$ satisfying the required conditions. □

The next lemma provides an actual construction of the required diagonal operator.

Lemma 3.8. *For any* $\alpha \in (0,1)$ *there is* $L = L(\alpha) > 0$ *with the following property. Let* $(\tau_k)_{k=1}^{\infty}$ *be an increasing non-negative sequence satisfying* $\sum_{k=1}^{\infty} \tau_k 2^{-k} < \infty$*, and let*

$$\widetilde{Z} := \sum_{k=0}^{\infty} \tau_{k+1} \xi^k,$$

where $\xi^k = (\xi_1^k, \xi_2^k, \dots, \xi_n^k)$ *and* ξ_i^k *(*$i \le n$*,* $k = 0, 1, \dots$*) are jointly independent* 0-1 *random variables with* $\mathbb{P}\{\xi_i^k = 1\} = 2^{-k}$*. Further, let* $\delta \in (0,1]$*. Then there is a random positive contraction* $\widetilde{D}$ *taking values in* $\mathcal{D}_n$ *such that*

$$\|\widetilde{D}\widetilde{Z}\|_1 \le \frac{L}{\delta} \mathbb{E}\|\widetilde{Z}\|_1 = \frac{Ln}{\delta} \sum_{k=0}^{\infty} \tau_{k+1} 2^{-k} \;\; \textit{everywhere on the probability space},$$

and $\mathbb{E}(\det \widetilde{D})^{\alpha-1} \le \exp(\delta)$*.*

Proof. Let $L \geq 2e$ be a number which we will determine later. Now, for each $k \geq 0$, define random variables

$$\nu_k := |\{i : \xi_i^k \neq 0\}|$$

and

$$\eta_k := \begin{cases} \left(\frac{\delta\nu_k}{L2^{-k}n}\right)^{\nu_k}, & \text{if } \delta\nu_k \geq L2^{-k}n; \\ 1, & \text{otherwise.} \end{cases}$$

As building blocks of the contraction $\widetilde{D}$, let us consider random diagonal matrices $D^{(k)}$ with

$$d_{jj}^{(k)} := \begin{cases} 1, & \text{if } \xi_j^k = 0; \\ \min\left(1, \frac{L2^{-k}n}{\delta\nu_k}\right), & \text{otherwise,} \end{cases} \qquad j = 1, 2, \ldots, n.$$

Then $\det D^{(k)} = {\eta_k}^{-1}$ and $\|D^{(k)}\xi^k\|_1 \leq \frac{L2^{-k}n}{\delta} = \frac{L}{\delta}\mathbb{E}\|\xi^k\|_1$ (deterministically). Note that $D^{(k)}$ acts as a dilation on the span of $\{e_i : \xi_i^k \neq 0\}$ provided that $\nu_k \geq \frac{L2^{-k}n}{\delta} = \frac{L}{\delta}\mathbb{E}\nu_k$, and as an isometry on the orthogonal complement. We construct the required contraction $\widetilde{D}$ as the product of contractions $D^{(k)}$ by setting $\widetilde{D} := \prod_{k=0}^{\infty} D^{(k)}$. Then

$$\|\widetilde{D}\widetilde{Z}\|_1 \leq \Big\|\sum_{k=0}^{\infty} \tau_{k+1} D^{(k)}\xi^k\Big\|_1 \leq \frac{Ln}{\delta}\sum_{k=0}^{\infty}\tau_{k+1}2^{-k} = \frac{L}{\delta}\mathbb{E}\|\widetilde{Z}\|_1.$$

Note that

$$\mathbb{E}\big(\det \widetilde{D}\big)^{\alpha-1} = \mathbb{E}\prod_{k=0}^{\infty}{\eta_k}^{1-\alpha} = \prod_{k=0}^{\infty}\mathbb{E}{\eta_k}^{1-\alpha}.$$

Next, for every $k \geq 0$ we have

$$\begin{aligned} \mathbb{E}{\eta_k}^{1-\alpha} &\leq 1 + \sum_{m=\lceil L2^{-k}n/\delta\rceil}^{\infty}\left(\frac{\delta m}{L2^{-k}n}\right)^{m-\alpha m}\mathbb{P}\{\nu_k = m\} \\ &\leq 1 + \sum_{m=\lceil L2^{-k}n/\delta\rceil}^{\infty}\left(\frac{e\delta}{L}\right)^{m}\left(\frac{L2^{-k}n}{\delta m}\right)^{\alpha m}. \end{aligned}$$

In particular, for all k such that $L2^{-k}n/\delta \geq 1$, using the relation $L \geq 2e$, we obtain

$$\mathbb{E}{\eta_k}^{1-\alpha} \leq 1 + 2\Big(\frac{e\delta}{L}\Big)^{\lceil L2^{-k}n/\delta\rceil},$$

and for all k satisfying $L2^{-k}n/\delta < 1$, we get

$$\mathbb{E}{\eta_k}^{1-\alpha} \leq 1 + 2\frac{e\delta}{L}\big(L2^{-k}n/\delta\big)^{\alpha}.$$

Now, let us choose $L = L(\alpha)$ sufficiently large so that both

$$\sum_{k:\, L2^{-k}n/\delta\geq 1} 2\Big(\frac{e\delta}{L}\Big)^{\lceil L2^{-k}n/\delta\rceil} \quad \text{and} \quad \sum_{k:\, L2^{-k}n/\delta<1} 2\frac{e\delta}{L}\big(L2^{-k}n/\delta\big)^{\alpha}$$

are less than $\delta/2$. Then, multiplying the estimates for $\mathbb{E}{\eta_k}^{1-\alpha}$, we get

$$\mathbb{E}\Big(\prod_{k=0}^{\infty}\eta_k\Big)^{1-\alpha} \leq \exp(\delta),$$

and the result follows. □

Proof of Proposition 3.4. Fix admissible α, δ and p. Without loss of generality, the distribution of the coordinates of the random vector X is continuous on the real line. Indeed, otherwise we can replace every coordinate x_i with $|x_i|+u_i$, where $u_1, u_2, \dots, u_n$ are jointly independent with $x_1, x_2, \dots, x_n$ and each u_i is uniformly distributed on $[0, \theta]$ for a very small parameter $\theta > 0$ chosen so that $\mathbb{E}(|x_i| + u_i)^p \approx \mathbb{E}|x_i|^p$. Then the random diagonal contraction D constructed for the new vector $X' := (|x_i| + u_i)_{i=1}^n$, will also satisfy the required properties with respect to X.

Set $Y := (|x_1|^p, |x_2|^p, \dots, |x_n|^p)$ and let $\widetilde{Y}, \widetilde{Z}$ be random vectors on a space $(\widetilde{\Omega}, \widetilde{\Sigma}, \widetilde{\mathbb{P}})$ constructed in Lemma 3.7 with respect to Y. By Lemma 3.8 and in view of the relation

$$\mathbb{E}\xi \geq \sum_{k=0}^{\infty} 2^{-k-1}\tau_k(\xi)$$

we can find a random positive contraction $\widetilde{D}$ on $\widetilde{\Omega}$ taking values in $\mathcal{D}_n$ such that for some $L = L(\alpha) > 0$ we have

$$\|\widetilde{D}\widetilde{Y}\|_1 \leq \|\widetilde{D}\widetilde{Z}\|_1 \leq \frac{L}{\delta}\mathbb{E}\|\widetilde{Z}\|_1 \leq \frac{4L}{\delta}\mathbb{E}\|\widetilde{Y}\|_1 \quad \text{everywhere on } \widetilde{\Omega}$$

and

$$\mathbb{E}(\det \widetilde{D})^{\alpha-1} \leq \exp(\delta).$$

In general, the operator $\widetilde{D}$ is not a function of $\widetilde{Y}$, which creates (purely technical) issues in defining corresponding operator on the original space $(\Omega, \Sigma, \mathbb{P})$. For completeness, let us describe an elementrary discretization argument resolving the problem:

Let $\{B_z\}$ be a partition of $\mathbb{R}^n_+$ into Borel subsets, indexed over $z = (z_1, z_2, \dots, z_n) \in$

$(\mathbb{Z} \cup \{-\infty\})^n$ and defined by

$$B_z := \big\{W \in \mathbb{R}_+^n : \, W_i \in (2^{z_i - 1}, 2^{z_i}] \;\text{ for all } i = 1, 2, \ldots, n\big\}$$

(we set $W_i = 0$ for $z_i = -\infty$). Further, for every z we let

$$\widetilde{\Omega}_z := \big\{\widetilde{\omega} \in \widetilde{\Omega} : \, \widetilde{Y}(\widetilde{\omega}) \in B_z\big\}$$

and $Q_z := \widetilde{D}(\widetilde{\Omega}_z) = \{M \in \mathcal{D}_n : \, M = \widetilde{D}(\widetilde{\omega}) \;\text{ for some } \widetilde{\omega} \in \widetilde{\Omega}_z\}$. First, consider all $z \in (\mathbb{Z} \cup \{-\infty\})^n$ such that $\widetilde{\Omega}_z$ is non-empty. For each $\widetilde{\Omega}_z$ we can choose an operator D_z such that $\det D_z \geq \det M$ for all $M \in Q_z$, and either $D_z \in Q_z$ or it is a limit point of the elements from Q_z (as a subspace of $\mathcal{D}_n \subset \mathbb{R}^{n^2}$). Of course, the choice of D_z does not have to be unique. Otherwise, if $\widetilde{\Omega}_z$ is empty then we set

$$D_z := \min\big(1, \frac{4L}{\delta \sum_{i=1}^n 2^{z_i}} \mathbb{E}\|\widetilde{Y}\|_1\big) \mathrm{Id}_n.$$

Finally, define a function $h : \mathbb{R}_+^n \to \mathcal{D}_n$ by setting $h(W) := D_z$ for all $W \in B_z$ and $z \in (\mathbb{Z} \cup \{-\infty\})^n$. Observe that h is Borel. Further, by the choice of D_z's, we have $\det h(\widetilde{Y}) \geq \det \widetilde{D}$ everywhere on $\widetilde{\Omega}$, whence $\mathbb{E}(\det h(\widetilde{Y}))^{\alpha-1} \leq \exp(\delta)$. Next, by the choice of sets B_z, we have $\|M(W)\|_1 \leq 2\|M'(W')\|_1$ for *any* two couples $(M, W), (M', W') \in Q_z \times B_z$. Together with the conditions on $\widetilde{D}$ and the definition of D_z's, this implies $\|D_z(W)\|_1 \leq 8L\mathbb{E}\|\widetilde{Y}\|_1/\delta$ for all $W \in B_z$, whence

$$\|h(W)\, W\|_1 \leq \frac{8L}{\delta} \mathbb{E}\|\widetilde{Y}\|_1 \quad \text{everywhere on } \mathbb{R}_+^n.$$

Now, taking $T := h(Y)$, we obtain a random diagonal contraction on $(\Omega, \Sigma, \mathbb{P})$ such that

$$\|T^{1/p} X\|_p^p = \|TY\|_1 \leq \frac{8L}{\delta} \mathbb{E}\|X\|_p^p \quad \text{everywhere on } \Omega$$

and $\mathbb{E}(\det T)^{\alpha-1} \leq \exp(\delta)$. Finally, setting $D := T^{1/p}$, we get the required operator. □

The above statement can be "tensorized". In what follows, we are interested only in the case $p = 2$ and $\alpha = 1/2$.

Proposition 3.9. *There is an absolute constant $C > 0$ with the following property. Let $A = (a_{ij})$ be an $n \times n$ random matrix satisfying* (1.4)*, and let $\delta \in (0, 1]$. Then there is a random positive contraction D taking values in $\mathcal{D}_n$ such that the Euclidean norms of the* *rows of AD are uniformly bounded by $\frac{C}{\sqrt{\delta}}\sqrt{n}$ everywhere on the probability space, and*

$$\mathbb{E}\det D^{-1} \le \exp(\delta n).$$

Proof. Indeed, for any $i = 1, 2, \ldots, n$, let D_i be the positive contraction defined with respect to the i-th row of A using Proposition 3.4 (with parameters $\alpha = 1/2$, $p = 2$), so that $D_1, D_2, \ldots, D_n$ are jointly independent. Then the product of these contractions $D := \prod_{i=1}^n D_i$ satisfies the required conditions. $\square$

Remark 3.10. It is not difficult to see that for any positive contraction $M \in \mathcal{D}_n$ there is an element $\widetilde{M} \in \mathcal{D}_n^2$ such that $\widetilde{M} \le \sqrt{2}M$ and $\det \widetilde{M}^{-1} \le \det M^{-2}$. Indeed, this follows easily from the fact that for any number $t \in (0, 1]$ there is $\widetilde{t} \in \{1\} \cup \left\{2^{-2^k}\right\}_{k=0}^{\infty}$ with $t^2 \le \widetilde{t} \le \sqrt{2}t$ (the constant $\sqrt{2}$ on the right-hand side is achieved for $t = \sqrt{2}/2 - o(1)$). Hence, the above statement implies that, given a matrix A satisfying (1.4) and a number $\delta > 0$, one can construct a random contraction $\widetilde{D}$ taking values in $\mathcal{D}_n^2$ such that each row of $A\widetilde{D}$ has Euclidean norm at most $\frac{C}{\sqrt{\delta}}\sqrt{n}$ (for some absolute constant $C > 0$), and $\mathbb{E}\det \widetilde{D}^{-1/2} \le \exp(\delta n)$.

3.2.3 Proof of the Theorem 3.2

The main result of the section is

Theorem 3.11. *Let $\delta \in (0, 1]$ and let $A = (a_{ij})$ be an $n \times n$ random matrix satisfying* (1.4). *Then*

$$\mathbb{P}\Big\{\exists D \in \mathcal{D}_n^2 : \det D \ge \exp(-\delta n) \textit{ and } \|AD\|_{\infty\to 2} \le \frac{C_{3.11}}{\sqrt{\delta}} n\Big\} \ge 1 - 4\exp(-\delta n/8),$$

where $C_{3.11} > 0$ is a universal constant.

Remark 3.12. The above theorem can be seen as a way to "regularize" the random matrix A by reducing its norm while preserving its "structure". In this connection, let us mention work [LLV17] where a very general problem of regularizing random matrices was discussed (see [LLV17, Section 5.4]).

As we have mentioned in the introduction, Theorem 3.2 follows almost immediately from the above statement; we give the proof of Theorem 3.2 at the very end of the section. The section is organized as follows. First, we use $\widetilde{D}$ constructed in Remark 3.10 to verify Theorem 3.11 under an additional assumption that the entries of A are symmetrically distributed (see Proposition 3.14). Then, we will apply a symmetrization procedure to prove Theorem 3.11 in full generality.

The next proposition implies that for a random matrix A satisfying (1.4) with symmetrically distributed entries and the operator $\widetilde{D}$ from Remark 3.10, the norm $\|A\widetilde{D}\|_{\infty\to 2}$ can be efficiently bounded from above as long as $\widetilde{D}$ is a Borel function of $|A|$ (here and further in the text, given a matrix $B = (b_{ij})$, by $|B|$ we shall denote the matrix $(|b_{ij}|)$).

Proposition 3.13. *Let $K > 0$ and let A be an $n \times n$ random matrix satisfying* (1.4)*, with symmetrically distributed entries. Further, let $\mathcal{F} \subset \mathcal{D}_n$ be any countable subset. Denote by $\mathcal{E}$ the event*

$$\mathcal{E} := \big\{\exists D \in \mathcal{F} : \text{ all rows of } AD \text{ have Euclidean norms at most } K\sqrt{n}\big\}.$$

Then

$$\mathbb{P}\{\exists D \in \mathcal{F} : \|AD\|_{\infty\to 2} \le CKn\} \ge \mathbb{P}(\mathcal{E}) - \exp(-n),$$

where $C > 0$ is an absolute constant.

Proof. Fix any admissible K and $\mathcal{F}$. Clearly, for any $n \times n$ matrix B and a diagonal matrix D, the Euclidean norms of rows of BD and $|B|D$ are the same. Hence, we may assume that there is a Borel function $f : \mathbb{R}_+^{n\times n} \to \mathcal{F}$ such that

$$\mathcal{E} = \big\{\text{all rows of } A\, f(|A|) \text{ have norms at most } K\sqrt{n}\big\}.$$

For any $D \in \mathcal{F}$, let

$$\mathcal{E}_D := \mathcal{E} \cap \{f(|A|) = D\}.$$

Without loss of generality, $\mathbb{P}(\mathcal{E}_D) > 0$ for any $D \in \mathcal{F}$.

Next, as the unit cube $[-1, 1]^n$ is the convex hull of its vertices $V = \{-1, 1\}^n$, we have

$$\|Af(|A|)\|_{\infty\to 2} = \sup_{y\in B_\infty^n} \|Af(|A|)y\| = \sup_{v\in V} \|Af(|A|)v\|. \tag{3.2}$$

Note that, given event $\mathcal{E}_D$, the entries of $Af(|A|) = AD$ are symmetrically distributed, so the distribution of ADv given $\mathcal{E}_D$ is the same for any vertex $v \in V$. Fix a vertex v.

Observe that for any $t > 0$ we have

$$\mathbb{P}_{\mathcal{E}_D}\{\|ADv\| > t\} \le \sup_B \mathbb{P}\{\|\widetilde{B}Dv\| > t\}, \tag{3.3}$$

where by $\mathbb{P}_{\mathcal{E}_D}$ we denote the conditional probability given $\mathcal{E}_D$ and the supremum is taken over all matrices $B = (b_{ij})$ such that the rows of BD have Euclidean norms at most $K\sqrt{n}$, and $\widetilde{B} = (r_{ij}b_{ij})$, with r_{ij} being jointly independent Rademacher (± 1) variables. Fix any admissible $B = (b_{ij})$.

Then the variables $\langle \widetilde{B}Dv, e_i\rangle$, $i = 1, 2, \dots, n$, are jointly independent and, in view of Lemma 2.5 and the choice of B, each variable $K^{-1}n^{-1/2}\langle \widetilde{B}Dv, e_i\rangle$ is $C_{2.5}$-subgaussian. By Lemma 2.6, there is an absolute constant $C > 0$ such that

$$\mathbb{P}\{\|\widetilde{B}Dv\| > CKn\} = \mathbb{P}\Big\{\frac{1}{n}\sum_{i=1}^{n}\langle \widetilde{B}Dv, e_i\rangle^2 > (CK)^2 n\Big\} \le \exp\big(-(1+\ln 2)n\big).$$

Then, taking a union bound over 2^n vertices of the unit cube and using (3.3) and (3.2), we get an estimate

$$\mathbb{P}_{\mathcal{E}_D}\{\|AD\|_{\infty\to 2} > CKn\} \le 2^n \cdot \sup_B \mathbb{P}\{\|\widetilde{B}Dv\| > CKn\} \le \exp(-n).$$

Finally, clearly

$$\mathbb{P}\{\|AD\|_{\infty\to 2} > CKn\} \le \mathbb{P}(\mathcal{E}^c) + \sum_D \mathbb{P}_{\mathcal{E}_D}\{\|AD\|_{\infty\to 2} > CKn\}\mathbb{P}(\mathcal{E}_D) \le \mathbb{P}(\mathcal{E}^c) + \exp(-n),$$

and the result follows. □

Proposition 3.14. *Let $\delta \in (0, 1]$ and let $A = (a_{ij})$ be an $n \times n$ random matrix satisfying* (1.4)*, with symmetrically distributed entries. Then*

$$\mathbb{P}\Big\{\exists D \in \mathcal{D}_n^2 : \det D \ge \exp(-\delta n) \textit{ and } \|AD\|_{\infty\to 2} \le \frac{C_{3.14}}{\sqrt{\delta}} n\Big\} \ge 1 - 2\exp(-\delta n/4).$$

Proof. Fix any $\delta \in (0, 1]$. In view of Remark 3.10, there is a random contraction D taking values in $\mathcal{D}_n^2$ such that each row of AD has the Euclidean norm at most $\frac{C}{\sqrt{\delta}}\sqrt{n}$ and $\mathbb{E}\det D^{-1/2} \le \exp(\delta n/4)$. Denote by $\mathcal{E}$ the event

$$\mathcal{E} := \big\{\det D \ge \exp(-\delta n)\big\}.$$

In view of the conditions on D and Markov's inequality, we have

$$\mathbb{P}(\mathcal{E}) \ge 1 - \exp(-\delta n/4).$$

Hence, by Proposition 3.13, taking $\mathcal{F}$ to be the set of all contractions from $\mathcal{D}_n^2$ having determinant at least $\exp(-\delta n)$, we obtain

$$\mathbb{P}\Big\{\exists D \in \mathcal{D}_n^2 : \det D \ge \exp(-\delta n) \text{ and } \|AD\|_{\infty\to 2} \le \frac{C_{3.14}}{\sqrt{\delta}} n\Big\} \ge 1 - \exp(-\delta n/4) - \exp(-n)$$

for a universal constant $C_{3.14} > 0$. □

The next statement shall be used in a symmetrization argument within the proof of Theorem 3.11; we think it may be of interest in itself.

Proposition 3.15. *Let $B = (b_{ij})$ be a non-random $n \times n$ matrix such that the Euclidean norm of every row is at most $\sqrt{n}$ and such that*

$$\Big|\sum_{j=1}^{n} b_{ij}\Big| \leq \sqrt{n}, \quad i = 1, 2, \dots, n.$$

Further, let π_i ($i = 1, 2, \dots, n$) be independent random permutations uniformly distributed on Π_n, and denote by $\widetilde{B} = (\widetilde{b}_{ij})$ the random $n \times n$ matrix with entries defined by

$$\widetilde{b}_{ij} := b_{i,\pi_i(j)}.$$

Then

$$\mathbb{P}\{\|\widetilde{B}\|_{\infty\to 2} \leq C_{3.15} n\} \geq 1 - \exp(-n)$$

for a universal constant $C_{3.15} > 0$.

Proof. We will show that for any $v \in \{-1, 1\}^n$ we have

$$\mathbb{P}\{\|\widetilde{B}v\| > C_{3.15} n\} \leq \exp(-n - n \ln 2)$$

for a sufficiently large universal constant $C_{3.15}$ and then take the union bound over the vertices of the cube.

Fix any $v = (v_1, v_2, \dots, v_n) \in \{-1, 1\}^n$ and let m be the number of ones in $(v_1, \dots, v_n)$. Clearly, the random variables $\langle \widetilde{B}v, e_i\rangle$ ($i = 1, 2, \dots, n$) are independent. Next, for a fixed i, the distribution of $\langle \widetilde{B}v, e_i\rangle$ coincides with that of the variable $\xi_i := \sum_{j=1}^{n} v_{\pi_i(j)} b_{ij}$. By Lemma 2.14 and in view of the condition on the rows of B, we have

$$\mathbb{P}\big\{|\xi_i - \mathbb{E}\xi_i| > \tau\big\} \leq 2\exp\Big(-\frac{\tau^2}{64n}\Big), \quad \tau > 0.$$

Hence, the variables $n^{-1/2}(\xi_i - \mathbb{E}\xi_i)$ ($i = 1, 2, \dots, n$) are C-subgaussian for an absoute constant $C > 0$. In view of Lemma 2.6, we get that

$$\mathbb{P}\Big\{\sum_{i=1}^{n} (\xi_i - \mathbb{E}\xi_i)^2 > \widetilde{C} n^2\Big\} \leq \exp(-n - n \ln 2) \tag{3.4}$$

for some constant $\widetilde{C} > 0$. Finally, observe that

$$\sum_{i=1}^{n} \xi_i^2 \leq 2\sum_{i=1}^{n} (\xi_i - \mathbb{E}\xi_i)^2 + 2\sum_{i=1}^{n} \big(\mathbb{E}\xi_i\big)^2 \quad \text{(deterministically)},$$

so, applying the estimate

$$\big|\mathbb{E}\xi_i\big| = \Big|\frac{2m-n}{n}\sum_{j=1}^{n} b_{ij}\Big| \leq \sqrt{n}$$

and (3.4), we obtain

$$\mathbb{P}\big\{\|\widetilde{B}v\|^2 > (2\widetilde{C}+2)n^2\big\} = \mathbb{P}\Big\{\sum_{i=1}^{n} \xi_i^2 > (2\widetilde{C}+2)n^2\Big\} \leq \exp(-n - n\ln 2).$$

□

Proof of Theorem 3.11. Let $\widetilde{A}$ be an independent copy of A. Obviously

$$\mathbb{E}\Big(\sum_{j=1}^{n} \widetilde{a}_{ij}\Big)^2 = \mathbb{E}\sum_{j=1}^{n} \widetilde{a}_{ij}^2 = n$$

for every $i = 1, 2, \ldots, n$. Then, in view of Markov's inequality, each row of $\widetilde{A}$ satisfies

$$\Big|\sum_{j=1}^{n} \widetilde{a}_{ij}\Big| \leq \sqrt{\frac{32n}{\delta}} \;\text{ and }\; \sum_{j=1}^{n} \widetilde{a}_{ij}^2 \leq \frac{32n}{\delta}$$

with probability at least $1 - \delta/16 > \exp(-\delta/8)$. Denote by $\widetilde{\mathcal{E}}$ the event

$$\widetilde{\mathcal{E}} := \Big\{\Big|\sum_{j=1}^{n} \widetilde{a}_{ij}\Big| \leq \sqrt{\frac{32n}{\delta}} \;\text{ and }\; \sum_{j=1}^{n} \widetilde{a}_{ij}^2 \leq \frac{32n}{\delta} \;\text{ for all } i = 1, 2, \ldots, n\Big\}.$$

In view of the above, $\mathbb{P}(\widetilde{\mathcal{E}}) \geq \exp(-\delta n/8)$. Let $\pi_1, \pi_2, \ldots, \pi_n$ be random permutations uniformly distributed on Π_n and jointly independent with $\widetilde{A}$, and denote by $\widetilde{B} = (\widetilde{b}_{ij})$ the random matrix with entries $\widetilde{b}_{ij} := \widetilde{a}_{i,\pi_i(j)}$ $(i, j \leq n)$. Then Proposition 3.15 yields

$$\mathbb{P}\Big\{\|\widetilde{B}\|_{\infty\to 2} \leq C_{3.15}\sqrt{n}\max_{i\leq n}\Big(\sum_{j=1}^{n} \widetilde{a}_{ij}^2\Big)^{1/2} \;\Big|\; \widetilde{A}\Big\} \geq 1 - \exp(-n),$$

whence, in particular,

$$\mathbb{P}\big\{\|\widetilde{B}\|_{\infty\to 2}\le C_{3.15}\sqrt{32/\delta}\,n\,|\,\widetilde{\mathcal{E}}\big\}\ge 1-\exp(-n).$$

But $\widetilde{B}$ is equidistributed with $\widetilde{A}$ given $\widetilde{\mathcal{E}}$, so that

$$\mathbb{P}\big\{\|\widetilde{A}\|_{\infty\to 2}\le C_{3.15}\sqrt{32/\delta}\,n\,|\,\widetilde{\mathcal{E}}\big\}\ge 1-\exp(-n).$$

Clearly, $\|\widetilde{A}D\|_{\infty\to 2}\le\|\widetilde{A}\|_{\infty\to 2}$ for any contraction $D\in\mathcal{D}_n$ (deterministically), so we obtain for the event $\mathcal{E}_1:=\big\{\|\widetilde{A}D\|_{\infty\to 2}\le C_{3.15}\sqrt{32/\delta}\,n\ \text{ for all } D\in\mathcal{D}_n\big\}$:

$$\mathbb{P}(\mathcal{E}_1)\ge(1-\exp(-n))\mathbb{P}(\widetilde{\mathcal{E}})\ge\frac{1}{2}\exp(-\delta n/8).$$

Next, the matrix $2^{-1/2}(A-\widetilde{A})$ has symmetrically distributed entries, and satisfies conditions of Proposition 3.14. Hence,

$$\begin{aligned}&\mathbb{P}\Big\{\|(A-\widetilde{A})D\|_{\infty\to 2}\le C_{3.14}\sqrt{2/\delta}\,n\ \text{ for some } D\in\mathcal{D}_n^2 \text{ with } \det D\ge\exp(-\delta n)\Big\}\\&\quad\ge 1-2\exp(-\delta n/4).\end{aligned}$$

Conditioning on $\mathcal{E}_1$, we get

$$\begin{aligned}&\mathbb{P}\Big\{\|(A-\widetilde{A})D\|_{\infty\to 2}\le C_{3.14}\sqrt{2/\delta}\,n\ \text{ for some } D\in\mathcal{D}_n^2 \text{ with } \det D\ge\exp(-\delta n)\,|\,\mathcal{E}_1\Big\}\\&\ge 1-\frac{2\exp(-\delta n/4)}{\mathbb{P}(\mathcal{E}_1)}\\&\ge 1-4\exp(-\delta n/8).\end{aligned}$$

Note that, given $\mathcal{E}_1$, we have $\|AD\|_{\infty\to 2}\le\|(A-\widetilde{A})D\|_{\infty\to 2}+C_{3.15}\sqrt{32/\delta}\,n$ *for all* contractions $D\in\mathcal{D}_n$. Combining this with the last formula, we obtain

$$\begin{aligned}&\mathbb{P}\Big\{\|AD\|_{\infty\to 2}\le C_{3.14}\sqrt{2/\delta}\,n+C_{3.15}\sqrt{32/\delta}\,n\\&\qquad\text{for some } D\in\mathcal{D}_n^2 \text{ with } \det D\ge\exp(-\delta n)\,|\,\mathcal{E}_1\Big\}\ge 1-4\exp(-\delta n/8).\end{aligned}$$

Finally, since A is independent from $\mathcal{E}_1$, the conditioning in the last estimate can be dropped, and we obtain the statement. □

To complete the proof of Theorem 3.2, we will need two more technical lemmas:

Lemma 3.16. *For any $\delta \in (0, 1/2]$ and all $n \in \mathbb{N}$ we have*

$$\left|\left\{D \in \mathcal{D}_n^2 : \det D \geq \exp(-\delta n)\right\}\right| \leq \left(\frac{2e}{\delta}\right)^{4\delta n}.$$

Proof. Denote $\mathcal{S} := \left\{D \in \mathcal{D}_n^2 : \det D \geq \exp(-\delta n)\right\}$. Note that for any matrix $D \in \mathcal{S}$ and for any $k \geq 0$, the number of diagonal elements of D equal to 2^{-2^k} is less than $2^{-k+1}\delta n$. Hence, the cardinality of $\mathcal{S}$ can be estimated as

$$|\mathcal{S}| \leq \prod_{k=0}^{\infty} \binom{n}{[2^{-k+1}\delta n]} \leq \prod_{k=0}^{\infty} \Big(\frac{e}{\delta}\Big)^{2^{-k+1}\delta n} 2^{k2^{-k+1}\delta n} = \Big(\frac{e}{\delta}\Big)^{4\delta n} 2^{4\delta n} = \Big(\frac{2e}{\delta}\Big)^{4\delta n}.$$

$\square$

Lemma 3.17. *For any $n \in \mathbb{N}$ and $K \in [2, 2\sqrt{n}]$, the unit Euclidean ball B_2^n can be covered by at most $(2eK^2)^{8n/K^2}$ translates of the cube $\frac{K}{\sqrt{n}}B_\infty^n$.*

Proof. First, note that for any $y \in B_2^n$ we have

$$\left|\left\{i \leq n : |y_i| \geq \frac{K}{2\sqrt{n}}\right\}\right| \leq \frac{4n}{K^2}.$$

Hence, it is sufficient to show that the set $|\{y \in B_2^n : |\mathrm{supp}(y)| \leq \frac{4n}{K^2}\}|$ can be covered by at most $(2eK^2)^{8n/K^2}$ translates of $\frac{K}{2\sqrt{n}}B_\infty^n$. A simple volumetric argument, together with an estimate $\mathrm{Vol}(B_2^m) \leq \left(\frac{2\pi e}{m}\right)^{m/2}$, implies that B_2^m can be covered by at most 7^m translates of $\frac{1}{\sqrt{m}}B_\infty^m$ (for any $m \in \mathbb{N}$). As a consequence, we obtain a covering of $B_2^{\lceil 4n/K^2 \rceil}$ by at most $7^{\lceil 4n/K^2 \rceil}$ translates of $\frac{K}{2\sqrt{n}}B_\infty^n$. Finally, the cardinality of the optimal covering of $|\{y \in B_2^n : |\mathrm{supp}(y)| \leq \frac{4n}{K^2}\}|$ can be estimated from above by

$$\binom{n}{\lceil 4n/K^2 \rceil} 7^{\lceil 4n/K^2 \rceil} \leq \left(2eK^2\right)^{8n/K^2}.$$

$\square$

Proof of Theorem 3.2. Let $\delta \in (0, 1/4]$ and $n \geq \frac{1}{4\delta}$. First, applying Lemma 3.17 with $K = 1/\sqrt{\delta}$, we see that B_2^n can be covered by $(2e/\delta)^{8n\delta}$ translates of the dilated cube $\frac{1}{\sqrt{n\delta}}B_\infty^n$. Let

$$\mathcal{Q} = \{D \in \mathcal{D}_n^2 : \det D \geq \exp(-\delta n)\}.$$

Then, in view of Lemma 3.16, we get that B_∞^n can be covered by at most $(2e/\delta)^{4\delta n}\exp(\delta n)$ parallelepipeds in such a way that for any $y \in B_\infty^n$ and $D \in \mathcal{Q}$, y is covered by a translate of $D(B_\infty^n)$. Combining the two coverings, we get a collection $\mathcal{C}$ of parallelepipeds covering

B_2^n such that

$$|\mathcal{C}| \leq (2e/\delta)^{4\delta n} \exp(\delta n) \cdot (2e/\delta)^{8n\delta} = \exp(\delta n + 12n\delta \ln \frac{2e}{\delta}),$$

and for any $y \in B_2^n$ and $D \in \mathcal{Q}$, the set $\mathcal{C}$ contains a translate of $D(\frac{1}{\sqrt{n\delta}} B_\infty^n)$ covering y. Finally, applying Theorem 3.11, we get that with probability at least $1 - 4\exp(-\delta n/8)$ for some $D \in \mathcal{Q}$ we have $AD(B_\infty^n) \subset \frac{C_{3.11} n}{\sqrt{\delta}} B_2^n$, implying

$$\mathbb{P}\Big\{\forall\ x \in B_2^n\ \ \exists\ P \in \mathcal{C}\ \text{ such that }\ x \in P \text{ and } A(P) \subset Ax + \frac{2 \cdot C_{3.11}\sqrt{n}}{\delta} B_2^n\Big\}$$
$$\geq 1 - 4\exp(-\delta n/8).$$

(the multiple "2" in the last formula appears because the translation $-Ax + A(P)$ is not origin-symmetric in general). □

Proof of 3.3. Fix n and δ, and let $\mathcal{C}$ be the collection of parallelepipeds defined in Theorem 3.2. For each $P \in \mathcal{C}$, choose a point $y_p \in P \cap B_2^n$, and let $\mathcal{N} := \{y_P :\ P \in \mathcal{C}\}$. Then, clearly,

$$|\mathcal{N}| = |\mathcal{C}| \leq \exp(\delta n + 12n\delta \ln \frac{2e}{\delta}),$$

and with probability at least $1 - 4\exp(-\delta n/8)$ for every $x \in B_2^n$ there is $y = y(x) \in \mathcal{N}$ with $-Ax + Ay \in \frac{C\sqrt{n}}{\delta} B_2^n$. In short,

$$\mathbb{P}\Big\{A(B_2^n) \subset \bigcup_{y \in A(\mathcal{N})} \big(y + \frac{C\sqrt{n}}{\delta} B_2^n\big)\Big\} \geq 1 - 4\exp(-\delta n/8).$$

□

3.3 The smallest singular value

Now, we proceed to the proof of the Theorem 3.1. As we already mentioned in Section 3.1, it heavily relies on results obtained by Rudelson and Vershynin in papers [RV08] and [RV09]. In the next section, we will state several intermediate results from those papers that we will need in Subsection 3.3.2 to complete our proof.

3.3.1 Background results

A crucial step in the proof of [RV08, Theorem 1.2] is a decomposition of the unit sphere into sets of "compressible" and "incompressible" vectors.

Definition 3.18 (Sparse, compressible and incompressible vectors)**.** Fix parameters $\theta, \rho \in (0,1)$. A vector $x \in \mathbb{R}^n$ is called *θn-sparse* if $|\mathrm{supp}(x)| \le \theta n$. A vector $x \in S^{n-1}$ is called *compressible* if x is within Euclidean distance ρ from the set of all θn-sparse vectors. Otherwise, x will be called *incompressible*. The set of all compressible unit vectors will be denoted by $\mathrm{Comp}_n(\theta, \rho)$, and the set of incompressible vectors — by $\mathrm{Incomp}_n(\theta, \rho)$. Sometimes, when the dimension n or the parameters θ, ρ are clear from the context, we will simply write Comp, Incomp to denote the sets.

Remark 3.19. A similar decomposition of the unit sphere was already introduced in an earlier paper [LPRTJ05] for the purpose of bounding the smallest singular value of rectangular matrices.

Obviously, for any $\varepsilon > 0$ we have

$$\mathbb{P}\{s_n(A) < \varepsilon n^{-1/2}\} \le \mathbb{P}\Big\{\inf_{y \in \mathrm{Comp}} \|Ay\| < \varepsilon n^{-1/2}\Big\} + \mathbb{P}\Big\{\inf_{y \in \mathrm{Incomp}} \|Ay\| < \varepsilon n^{-1/2}\Big\}.$$

Treatment of the compressible vectors is simpler due to the fact the the set Comp is "small"; we will deal with this set in the first part of Section 3.3.2. Let us remark that, unlike in the subgaussian result of [RV08], where an estimate for compressible vectors follows almost directly from an analogue of Lemma 3.26 (see below) together with a standard covering argument, in our case we will still need to use additional results (proved in Section 3.2.3) as the norm $\|A\|$ may be "too large". We will need the following simple lemma:

Lemma 3.20. *For any $\theta, \rho \in (0,1]$ the set $\mathrm{Comp} = \mathrm{Comp}_n(\theta, \rho)$ admits a Euclidean 3ρ-net $\mathcal{N} \subset \mathrm{Comp}$ of cardinality $|\mathcal{N}| \le (e/\theta)^{\theta n} \big(\frac{5}{\rho}\big)^{\theta n}$.*

Proof. Note that the definition of Comp implies that for any $y \in \mathrm{Comp}$ there is $y' \in S^{n-1}$ such that $|\mathrm{supp}(y')| \le \theta n$ and $\|y - y'\| \le 2\rho$. Hence, it is enough to show that one can find a Euclidean ρ-net $\mathcal{N}$ on the set of θn-sparse unit vectors, with the required estimate on $|\mathcal{N}|$. This follows from a standard estimate on the cardinality of an optimal ρ-net on $S^{\lfloor \theta n \rfloor - 1}$, together with a bound for the binomial coefficient $\binom{n}{\lfloor \theta \rfloor}$. $\square$

Incompressible vectors have the important property that a significant portion of their coordinates are of order $n^{-1/2}$. In paper [RV08], this property was referred to as "incompressible vectors are spread". For reader's convenience, we provide a proof of this fact below (let us note once again that analogous concepts were already considered in [LPRTJ05]).

Lemma 3.21 ([RV08, Lemma 3.4])**.** *For any $\theta, \rho \in (0,1)$ and for any vector $x \in \mathrm{Incomp}_n(\theta, \rho)$ there is a subset of indices $\sigma(x) \subset \{1, 2, \ldots, n\}$ of cardinality at least $\frac{1}{2}\rho^2\theta n$ such that for* *all $i \in \sigma(x)$ we have*

$$\frac{\rho}{\sqrt{2n}} \le x_i \le \frac{1}{\sqrt{\theta n}}.$$

Proof. For every subset $I \subset \{1, 2, \ldots, n\}$, let P_I be the coordinate projection onto the span of $\{e_i : i \in I\}$. Let $\sigma = \sigma(x) := \sigma_1 \cap \sigma_2$, where

$$\sigma_1 = \Big\{ i \le n : |x_i| \le \frac{1}{\sqrt{\theta n}} \Big\}, \qquad \sigma_2 = \Big\{ i \le n : |x_i| \ge \frac{\rho}{\sqrt{2n}} \Big\}.$$

Since $\|x\| = 1$, we have $|\sigma_1^c| \le \theta n$, and $P_{\sigma_1^c}(x)$ is a θn-sparse vector. Then the condition that x is incompressible implies $\|P_{\sigma_1}(x)\| = \|x - P_{\sigma_1^c}(x)\| > \rho$. Hence,

$$\|P_\sigma(x)\|^2 \ge \|P_{\sigma_1}(x)\|^2 - \|P_{\sigma_2^c}(x)\|^2 \ge \rho^2 - n \cdot \|P_{\sigma_2^c}(x)\|_\infty^2 \ge \rho^2/2. \tag{3.5}$$

On the other hand, in view of the inclusion $\sigma(x) \subset \sigma_1$, we get

$$\|P_\sigma(x)\|^2 \le \|P_\sigma(x)\|_\infty^2 \cdot |\sigma| \le \frac{1}{\theta n} \cdot |\sigma|. \tag{3.6}$$

Together (3.5) and (3.6) imply that $|\sigma| \ge \frac{1}{2}\rho^2 \theta n$. □

For incompressible vectors we will need the following basic estimate from [RV08].

Proposition 3.22 ([RV08, Lemma 3.5]). *Let M be a random $n \times n$ matrix with column vectors $X^1, X^2, \ldots, X^n$, and let H_j ($j = 1, 2, \ldots, n$) be the span of all column vectors except the j-th. Then for every $\varepsilon > 0$ we have*

$$\mathbb{P}\Big\{ \inf_{y \in \mathrm{Incomp}(\theta,\rho)} \|My\| < \varepsilon \rho n^{-1/2} \Big\} \le \frac{1}{\theta n} \sum_{j=1}^n \mathbb{P}\{ \mathrm{dist}(X^j, H_j) < \varepsilon \}.$$

In view of independence and equi-measurability of the columns of A in our model, the above proposition yields for any $\varepsilon > 0$:

$$\mathbb{P}\Big\{ \inf_{y \in \mathrm{Incomp}(\theta,\rho)} \|Ay\| < \varepsilon \rho n^{-1/2} \Big\} \le \frac{1}{\theta} \mathbb{P}\Big\{ \Big| \sum_{i=1}^n X_i^* A_{in} \Big| < \varepsilon \Big\},$$

where $X^* = (X_1^*, X_2^*, \ldots, X_n^*)$ denotes a random normal unit vector to the span of the first $n-1$ columns of A. Obtaining small ball probability estimates for $\Big| \sum_{i=1}^n X_i^* A_{in} \Big|$ was a crucial ingredient of [RV08].

Given a real-valued random variable ξ, define its *Levy concentration function* as

$$\mathcal{L}(\xi, z) := \sup_{\lambda \in \mathbb{R}} \mathbb{P}\{|\xi - \lambda| \le z\}, \;\; z \ge 0.$$

First, let us look at some well known estimates of $\mathcal{L}(\xi, v)$ and then state a stronger bound from [RV08].

Theorem 3.23 (Rogozin, [Rog61])**.** *Let $n \in \mathbb{N}$, let $\xi_1, \xi_2, \ldots, \xi_n$ be jointly independent random variables and let $t_1, t_2, \ldots, t_n$ be some positive real numbers. Then for any $t \ge \max_j t_j$ we have*

$$\mathcal{L}\Big(\sum_{j=1}^{n} \xi_j, t\Big) \le C_{3.23}\, t \left(\sum_{j=1}^{n} (1 - \mathcal{L}(\xi_j)) {t_j}^2\right)^{-1/2},$$

where $C_{3.23} > 0$ is an absolute constant.

Obviously, if ξ is essentially non-constant, there are $v > 0$ and $u \in (0,1)$ such that $\mathcal{L}(\xi, v) \le u$. The following lemma is an elementary consequence of Theorem 3.23 (see [LPRTJ05, Lemma 3.6] and [RV08, Lemma 2.6] for similar statements proved under additional moment assumptions on the variable).

Lemma 3.24. *Let ξ be a random variable with $\mathcal{L}(\xi, \widetilde{v}) \le \widetilde{u}$ for some $\widetilde{v} > 0$ and $\widetilde{u} \in (0,1)$. Then there are $v' > 0$ and $u' \in (0,1)$ depending only on $\widetilde{u}, \widetilde{v}$ with the following property: Let $\xi_1, \xi_2, \ldots, \xi_n$ be independent copies of ξ. Then for any vector $y \in S^{n-1}$ we have*

$$\mathcal{L}\Big(\sum_{j=1}^{n} y_j \xi_j, v'\Big) \le u'.$$

Proof. By Theorem 3.23, for any $y \in S^{n-1}$ and any $h \ge \max_j |y_j| \widetilde{v}$, we have

$$\mathcal{L}\Big(\sum_{j=1}^{n} y_j \xi_j, h\Big) \le \frac{C_{3.23} h}{\widetilde{v}\sqrt{1 - \widetilde{u}}}.$$

Define $v' := \frac{\widetilde{v}\sqrt{1-\widetilde{u}}}{2C_{3.23}}$ and consider two cases.
1) For every $j = 1, \ldots, n$ we have $|y_j| \le \frac{\sqrt{1-\widetilde{u}}}{2C_{3.23}}$. Then $v' \ge \max_j |y_j| \widetilde{v}$, and we obtain from the above relation

$$\mathcal{L}\Big(\sum_{j=1}^{n} y_j \xi_j, v'\Big) \le \frac{1}{2}.$$

2) There is j_0 such that $|y_{j_0}| > \frac{\sqrt{1-\widetilde{u}}}{2C_{3.23}}$. Then we get

$$\mathcal{L}\Big(\sum_{j=1}^{n} y_j\xi_j, v'\Big) \leq \mathcal{L}(y_{j_0}\xi_{j_0}, v') \leq \widetilde{u}.$$

Thus, we can take $u' := \max(1/2, \widetilde{u})$. □

Lemma 3.25 ("Tensorization lemma", [RV08, Lemma 2.2])**.** *Let $\alpha_1, \alpha_2, \ldots, \alpha_n$ be i.i.d. random variables, and let $\varepsilon_0 > 0$.*

- *Assume that*

  $$\mathcal{L}(\alpha_1, \varepsilon) \leq L\varepsilon \qquad \textit{for some } L > 0 \textit{ and for all } \varepsilon \geq \varepsilon_0.$$

  Then

  $$\mathbb{P}\Big\{\sum_{j=1}^{n} {\alpha_j}^2 \leq \varepsilon^2 n\Big\} \leq (CL\varepsilon)^n \qquad \textit{for all } \varepsilon \geq \varepsilon_0,$$

  where $C > 0$ is an absolute constant.

- *Assume that $\mathcal{L}(\alpha_1, v') \leq u'$ for some $v' > 0$ and $u' \in (0, 1)$. Then there are $v > 0$ and $u \in (0, 1)$ depending only on u', v' such that*

  $$\mathbb{P}\Big\{\sum_{j=1}^{n} {\alpha_j}^2 \leq vn\Big\} \leq u^n.$$

As a consequence of Lemmas 3.24 and 3.25, we get

Lemma 3.26. *Let α be a random variable with $\mathcal{L}(\alpha, \widetilde{v}) \leq \widetilde{u}$ for some $\widetilde{v} > 0$ and $\widetilde{u} \in (0, 1)$. Then there are $v > 0$ and $u \in (0, 1)$ depending only on $\widetilde{u}, \widetilde{v}$ with the following property: Let A be an $n \times n$ random matrix with i.i.d. entries equidistributed with α. Then for any $y \in S^{n-1}$ we have*

$$\mathbb{P}\big\{\|Ay\| \leq v\sqrt{n}\big\} \leq u^n.$$

Remark 3.27. Lemma 3.26 can be compared with [LPRTJ05, Proposition 3.4] and [RV08, Corollary 2.7]; however, those statements were proved with additional assumptions on the entries of A.

To get a stronger estimate than the one obtained in Lemma 3.24, the following notion was developed in [RV08] and [RV09] (see also preceding work [TV09] by Tao and Vu).

Definition 3.28 (Essential least common denominator)**.** For parameters $r \in (0,1)$ and $h > 0$ and any non-zero vector $x \in \mathbb{R}^n$, define

$$\mathrm{LCD}_{h,r}(x) := \inf\bigl\{t > 0 : \mathrm{dist}(tx, \mathbb{Z}^n) < \min(r\|tx\|, h)\bigr\}.$$

We note that later we shall choose r sufficiently small and h to be a small multiple of $\sqrt{n}$. Thus, most of the coordinates of $\mathrm{LCD}_{h,r}(x) \cdot x$ are within a small distance to integers. For a detailed discussion of the above notion, we refer to [Rud13].

The next statement is proved in [RV09].

Theorem 3.29 ([RV09, Theorem 3.4])**.** *Let $\xi_1, \xi_2, \dots, \xi_n$ be independent copies of a centered random variable such that $\mathcal{L}(\xi_i, v) \le u$ for some $v > 0$ and $u \in (0,1)$. Further, let $x = (x_1, x_2, \dots, x_n) \in S^{n-1}$ be a fixed vector. Then for every $h > 0$, $r \in (0,1)$ and for every*

$$\varepsilon \ge \frac{1}{\mathrm{LCD}_{h,r}(x)},$$

we have

$$\mathcal{L}\Bigl(\sum_{i=1}^n x_i\xi_i, \varepsilon v\Bigr) \le \frac{C_{3.29}\varepsilon}{r\sqrt{1-u}} + C_{3.29}\exp\bigl(-2(1-u)h^2\bigr),$$

where $C_{3.29}$ is a universal constant.

Thus, in order to get a satisfactory small ball probability estimate for the infimum over incompressible vectors, it is sufficient to show that the random normal X^* has exponentially large LCD with probability close to one. This will be done in the second part of Section 3.3.2. As for the set Comp, our treatment of the random normal will be based on results of Section 3.2.3.

3.3.2 Proof of the Theorem 3.1

In this section we give a proof of Theorem 3.1 stated in the introduction. Let us start with a version of Theorem 3.2 more convenient for us:

Theorem 3.30. *Let $\delta \in (0, 1/4]$, $n \ge \frac{1}{4\delta}$, $\varepsilon \in (0, 1/2]$, $S \subset S^{n-1}$, and let $\mathcal{N} \subset S$ be a Euclidean ε-net on S. Then there exists a (deterministic) subset $\widetilde{\mathcal{N}} \subset S$ with*

$$|\widetilde{\mathcal{N}}| \ \le \ \exp\bigl(13\delta n \ln \frac{2e}{\delta}\bigr)|\mathcal{N}|$$

such that for any $n \times n$ random matrix A satisfying (1.4)*, with probability at least $1 - 4\exp(-\delta n/8)$ the set $\widetilde{\mathcal{N}}$ is a $(\frac{\varepsilon C_\star}{\delta}\sqrt{n})$–net on S with respect to the pseudometric $d(x,y) := \|A(x-y)\|$ $(x, y \in S^{n-1})$, where $C_\star > 0$ is an absolute constant.*

Proof. Fix parameters n and δ, and let $\mathcal{C}$ be the collection of parallelepipeds from Theorem 3.2 covering B_2^n. Define a set $\widetilde{\mathcal{C}} := \big\{\varepsilon P+y:\ P\in\mathcal{C},\ y\in\mathcal{N},\ S\cap(\varepsilon P+y)\neq\emptyset\big\}$ and for every $\widetilde{P}\in\mathcal{C}$ let $y_{\widetilde{P}}$ be a point in the intersection $S\cap\widetilde{P}$. Finally, set $\widetilde{\mathcal{N}} := \{y_{\widetilde{P}}:\ \widetilde{P}\in\widetilde{\mathcal{C}}\}$. Informally speaking, $\widetilde{\mathcal{C}}$ is a "product" of the rescaled collection $\varepsilon\cdot\mathcal{C}$ and the net $\mathcal{N}$. For each parallelepiped in $\widetilde{\mathcal{C}}$ having a non-empty intersection with S, we take one (arbitrary) point from this intersection to construct the refined net $\widetilde{\mathcal{N}}$. What remains is to check that with high probability $\widetilde{\mathcal{N}}$ is indeed a $(\frac{\varepsilon C}{\delta}\sqrt{n})$–net on S with respect to the pseudometric $d(x,y) := \|A(x-y)\|$.

Observe that

$$|\widetilde{\mathcal{N}}| = |\widetilde{\mathcal{C}}| \le |\mathcal{C}|\cdot|\mathcal{N}| \le \exp\big(13\delta n\ln\frac{2e}{\delta}\big)|\mathcal{N}|.$$

Next, let A be an $n\times n$ random matrix satisfying (1.4), and define event $\mathcal{E}$ as

$$\mathcal{E} := \Big\{\forall\, x\in B_2^n \quad \exists P\in\mathcal{C} \text{ such that } x\in P \text{ and } A(P)\subset Ax+\frac{C\sqrt{n}}{\delta}B_2^n\Big\}.$$

By Theorem 3.2, we have $\mathbb{P}(\mathcal{E})\ge 1-4\exp(-\delta n/8)$.

Fix any point $x\in S$. By the definition of $\mathcal{N}$, there is a vector $y\in\mathcal{N}$ such that $\varepsilon^{-1}(x-y)\in B_2^n$. Hence, for any point $\omega\in\mathcal{E}$ on the probability space, there is a parallelepiped $P=P(\omega)\in\mathcal{C}$ such that $\varepsilon^{-1}(x-y)\in P$ and

$$A_\omega(P)\subset A_\omega\big(\varepsilon^{-1}(x-y)\big)+\frac{C\sqrt{n}}{\delta}B_2^n.$$

Note that $S\cap(\varepsilon P+y)\supset\{x\}\neq\emptyset$, whence $\widetilde{P}:=\varepsilon P+y\in\widetilde{\mathcal{C}}$, and, from the above relation,

$$A_\omega(\widetilde{P})\subset A_\omega x+\frac{\varepsilon C\sqrt{n}}{\delta}B_2^n,$$

whence

$$A_\omega y_{\widetilde{P}}-A_\omega x\subset\frac{\varepsilon C\sqrt{n}}{\delta}B_2^n,$$

where $y_{\widetilde{P}}\in\widetilde{\mathcal{N}}$. We have shown that

$$\mathcal{E}\subset\Big\{\forall x\in S\ \exists y=y(x)\in\widetilde{\mathcal{N}} \text{ such that } \|A(x-y)\|\le\frac{\varepsilon C\sqrt{n}}{\delta}\Big\},$$

and the result follows. $\square$

Remark 3.31. Let us note that a weaker version of Theorem 3.30, with condition $\widetilde{\mathcal{N}}\subset S$ dropped, can be proved by applying Corollary 3.3 instead of Theorem 3.2.

At this point, a significant part of our argument follows the same scheme as in [RV08].

In the first part of this section, we are dealing with compressible vectors.

Proposition 3.32 (Compressible vectors)**.** *Let α be a centered random variable with unit variance such that $\mathcal{L}(\alpha, \widetilde{v}) \leq \widetilde{u}$ for some $\widetilde{v} > 0$ and $\widetilde{u} \in (0,1)$. Then there are numbers $\theta_{3.32}, v_{3.32} > 0$ and $u_{3.32} \in (0,1)$ depending only on $\widetilde{v}, \widetilde{u}$ with the following property: Let $n \in \mathbb{N}$ and let A be an $n \times n$ random matrix with i.i.d. entries equidistributed with α. Then for* $\mathrm{Comp} = \mathrm{Comp}_n(\theta_{3.32}, \theta_{3.32})$ *we have*

$$\mathbb{P}\Big\{\inf_{y \in \mathrm{Comp}} \|Ay\| < v_{3.32}\sqrt{n}\Big\} \leq 5\, {u_{3.32}}^n.$$

Proof. Without loss of generality, we can assume that n is large. First, note that by Lemma 3.26 we have a strong probability estimate for any fixed unit vector: there are $v > 0$ and $u \in (0,1)$ depending on $\widetilde{v}, \widetilde{u}$ such that for any $y \in S^{n-1}$ we get

$$\mathbb{P}\{\|Ay\| < v\sqrt{n}\} \leq u^n. \tag{3.7}$$

In order to obtain a uniform estimate over a set $S = \mathrm{Comp}_n(\theta, \theta)$ for some small parameter θ, we will take a net $\mathcal{N} \subset S$ constructed in Lemma 3.20 and refine it with the help of Theorem 3.30 to get a net $\widetilde{\mathcal{N}}$ with respect to pseudometric $\|A(x-y)\|$. We will apply Theorem 3.30 with parameter δ defined as the largest number in $(0, 1/4]$ so that $\exp\big(13\delta n \ln \frac{2e}{\delta}\big) \leq u^{-n/3}$. Let us describe the procedure in more detail.

First, define parameter $\theta \in (0, 1/6]$ as the largest number satisfying the inequalities

$$\left(\frac{5e}{\theta^2}\right)^{\theta n} \leq u^{-n/3} \quad \text{and} \quad \frac{3\theta C_\star}{\delta} \leq \frac{v}{2}.$$

Let S be as above. By Lemma 3.20, there is a 3θ-net $\mathcal{N} \subset S$ on S (with respect to the usual Euclidean metric) of cardinality $|\mathcal{N}| \leq (\frac{5e}{\theta^2})^{\theta n}$. Now, by Theorem 3.30, there is a *deterministic* subset $\widetilde{\mathcal{N}} \subset S$ having the following properties:

- $|\widetilde{\mathcal{N}}| \leq \exp\big(13\delta n \ln \frac{2e}{\delta}\big) \cdot |\mathcal{N}| \leq u^{-n/3} \cdot \big(\frac{5e}{\theta^2}\big)^{\theta n} \leq u^{-2n/3}$;
- With probability at least $1 - 4\exp(-\delta n/8)$ for every $y \in S$ there exists $x(y) \in \widetilde{\mathcal{N}}$ such that
$$\|A(x-y)\| \leq \frac{3\theta \cdot C_\star}{\delta}\sqrt{n} \leq \frac{v}{2}\sqrt{n}.$$

Applying the union bound over $\widetilde{\mathcal{N}}$ to relation (3.7), we get

$$\mathbb{P}\{\|Ay'\| < v\sqrt{n} \text{ for some } y' \in \widetilde{\mathcal{N}}\} \leq |\widetilde{\mathcal{N}}| u^n \leq u^{n/3}.$$

On the other hand, the second property of $\widetilde{\mathcal{N}}$ implies that

$$\mathbb{P}\Big\{\inf_{y\in S}\|Ay\| < \inf_{y\in\widetilde{\mathcal{N}}}\|Ay\| - \frac{v\sqrt{n}}{2}\Big\} \le 4\exp(-\delta n/8).$$

Combining the two estimates, we get

$$\mathbb{P}\{\|Ay\| < v\sqrt{n}/2 \text{ for some } y\in S\} \le u^{n/3} + 4\exp(-\delta n/8),$$

and the result follows with $u_{3.32} := \max\{u^{1/3}, \exp(-\delta/8)\}$. □

Remark 3.33. It is not difficult to see that Proposition 3.32 can be stated and proved in the same way for A which is not square, but instead is an $n-1\times n$ matrix with i.i.d. entries equidistributed with α. Indeed, for n large enough we can assume that $\gamma\cdot n < (n-1) < n$ for γ as close to one as we want (the values of $\theta_{3.32}$, $u_{3.32}$ and $v_{3.32}$ may differ in that case). This will be important for us later.

Remark 3.34. Proposition 3.32 could be proved by a completely different argument based on [Tik15, Proposition 13] and not using results of Section 3.2.3 at all. However, we prefer to have a "uniform" treatment of both compressible and incompressible vectors.

Let us turn to estimating the infimum over incompressible vectors. As we already discussed in Section 3.3.1, it suffices to show that the random unit normal vector to the span of the first $n-1$ columns of A has exponentially large LCD with probability very close to one. This property is verified in Theorem 3.39 below. We start with some auxiliary statements. First, note that Theorem 3.29 together with Lemma 3.25 imply that anti-concentration probability for a single vector can be estimated in terms of the LCD of the vector. Namely, the bigger $\mathrm{LCD}(x)$ is, the less is the probability that the image Ax concentrates in a small ball:

Lemma 3.35 (Small ball probability for a single vector; see [RV08, Lemma 5.5])**.** *Let $h>0$, $r\in(0,1)$ and let α be a random variable satisfying $\mathcal{L}(\alpha,\widetilde{v}) \le \widetilde{u}$ for some $\widetilde{v}>0$ and $\widetilde{u}\in(0,1)$. Then there is $L_{3.35}\ge 1$ depending only on $\widetilde{v},\widetilde{u}$ with the following property: Let A' be an $n-1\times n$ random matrix with i.i.d. elements equidistributed with α. Then for any vector $x\in S^{n-1}$ and any*

$$\varepsilon \ge \widetilde{v}\cdot\max\Big(\frac{1}{\mathrm{LCD}_{h,r}(x)}, \exp(-2(1-\widetilde{u})h^2)\Big)$$

we have

$$\mathbb{P}\{\|A'x\| < \varepsilon\sqrt{n}\} \le (L_{3.35}\varepsilon/r)^{n-1}.$$

Proof. Fix any vector $x \in S^{n-1}$ and denote $Y = (Y_1, Y_2, \dots, Y_{n-1}) := A'x$. Note that, in view of Theorem 3.29, we have

$$\mathcal{L}\big(Y_i, \varepsilon\big) \le \frac{C_{3.29}\varepsilon}{r\sqrt{1-\widetilde{u}}} + C_{3.29}\exp\big(-2(1-\widetilde{u})h^2\big) \le \frac{C_{3.29}(1+\widetilde{v}^{-1})\varepsilon}{r\sqrt{1-\widetilde{u}}}, \qquad i \le n,$$

for any ε satisfying conditions of the lemma. Hence, by Lemma 3.25,

$$\mathbb{P}\Big\{\sum_{i=1}^{n-1} Y_i^2 \le \varepsilon^2(n-1)\Big\} \le \Big(\frac{C'(1+\widetilde{v}^{-1})\varepsilon}{r\sqrt{1-\widetilde{u}}}\Big)^{n-1}.$$

□

The above statement is useful for incompressible vectors: the following Lemma 3.36 shows that incompressible vectors have LCD at least of order $\sqrt{n}$. The lemma is taken from papers [RV08, RV09], and its proof is included for completeness.

Lemma 3.36 (see [RV09, Lemma 3.6])**.** *For every $\theta, \rho \in (0,1)$ there are $q_{3.36} = q_{3.36}(\theta, \rho) > 0$ and $r_{3.36} = r_{3.36}(\theta, \rho) > 0$ such that for every $h > 0$ any vector $x \in \mathrm{Incomp}_n(\theta, \rho)$ satisfies*

$$\mathrm{LCD}_{h, r_{3.36}}(x) \ge q_{3.36}\sqrt{n}.$$

Proof. Set $a := \frac{1}{2}\rho^2\theta$ and $b := \rho/\sqrt{2}$. We choose $r = r_{3.36} := b\sqrt{\frac{a}{2}} = \frac{1}{2}\rho^2\sqrt{\theta}$ and $q = q_{3.36} := \big(1/\sqrt{\theta} + \frac{2r}{a}\big)^{-1} = \sqrt{\theta}/3$.

Let $x \in \mathrm{Incomp}_n(\theta, \rho)$, $h > 0$ and assume that $\mathrm{LCD}_{h,r}(x) < q\sqrt{n}$. Then, by definition of least common denominator, there exist $p \in \mathbb{Z}^n$ and $\lambda \in (0, q\sqrt{n})$ such that

$$\|\lambda x - p\| < r\lambda < rq\sqrt{n} = \frac{1}{6}\rho^2\theta\sqrt{n} = \frac{1}{3}a\sqrt{n}. \tag{3.8}$$

It is easy to check that for a vector with such norm the set

$$\widetilde{\sigma}(x) := \big\{i \le n : |\lambda x_i - p_i| < 2/3\big\}$$

has a cardinality at least $(1-\frac{a^2}{4})n$. Further, by Lemma 3.21, the set of "spread" coordinates $\sigma(x)$ has cardinality at least an. Hence, the set $I(x) := \sigma(x) \cap \widetilde{\sigma}(x)$ is non-empty, and $|I(x)| > \frac{a}{2}n$. For any $i \in I(x)$ we have

$$|p_i| < \lambda|x_i| + \frac{2}{3} < \frac{q}{\sqrt{\theta}} + \frac{2rq}{a} = 1$$

(in the last step we used our definition of q). Since $p \in \mathbb{Z}^n$, we get that $p_i = 0$ for all

$i \in I(x)$.

Finally, due to the definition of $I(x)$ and our choice of r, denoting by P_J the coordinate projection on a span $\{i \in J : e_i\}$, we obtain

$$\|\lambda x - p\|^2 \geq \|\lambda P_I(x)\|^2 > \lambda^2 |I(x)| \frac{\rho^2}{2n} = \lambda^2 \frac{\rho^2 a}{4} = (r\lambda)^2,$$

which contradicts (3.8) and, hence, the assumption that $\mathrm{LCD}_{h,r}(x) < q\sqrt{n}$. □

Let $n \in \mathbb{N}$, $h > 0$, $\theta, \rho \in (0,1)$, and let $q_{3.36}$ and $r_{3.36}$ be as in the above statement. Following [RV08], we consider the "level sets" S_k of $\mathrm{Incomp}_n(\theta, \rho)$ defined as

$$S_k = S_k(\theta, \rho, h) := \big\{x \in \mathrm{Incomp}_n(\theta, \rho) \,:\, k \leq \mathrm{LCD}_{h, r_{3.36}}(x) < 2k\big\}, \quad k \geq 0.$$

In the proof of the theorem below we will partition $\mathrm{Incomp}_n(\theta, \rho)$ into subsets of vectors having LCD's of the same order:

$$\mathrm{Incomp}_n(\theta, \rho) = \bigsqcup_{k = 2^i, i \geq i_0} S_k, \tag{3.9}$$

where, using Lemma 3.36, we introduce the lower bound $i_0 := \log_2(q_{3.36}\sqrt{n}/2)$ (we have $S_k = \emptyset$ for all $k < q_{3.36}\sqrt{n}/2$). Following [RV08], we are going to combine estimates for individual sets S_k.

A principal observation made in [RV09] and [RV08] is that the sets S_k admit Euclidean ε-nets of relatively small cardinality. We give both the formal statement and its proof from [RV09] below for the sake of completeness:

Lemma 3.37 ([RV09, Lemma 4.8])**.** *For any $\theta, \rho \in (0,1)$ there is $L = L(\theta, \rho) > 0$ such that for every $h \geq 1$ and $k > 0$ the set S_k admits a Euclidean $(4h/k)$-net of cardinality at most $\big(kL/\sqrt{n}\big)^n$.*

Proof. In view of Lemma 3.36, we can assume that $k \geq q_{3.36}\sqrt{n}/2$. Further, without loss of generality $\frac{4h}{k} < 2$; otherwise a one-point net works.

Fix for a moment a point $x \in S_k$. Then, by definition of the "level sets", $k \leq \mathrm{LCD}_{h, r_{3.36}}(x) < 2k$. By definition of LCD, there exists $p = p(x) \in \mathbb{Z}^n$ such that

$$\|\,\mathrm{LCD}_{h, r_{3.36}}(x) \cdot x - p\| \leq h.$$

Hence,

$$\Big\|x - \frac{p}{\mathrm{LCD}_{h, r_{3.36}}(x)}\Big\| \leq \frac{h}{k} < \frac{1}{2}.$$

It is a simple planimetric observation that if we normalize the vector $p/\operatorname{LCD}_{h,r_{3.36}}(x)$, the distance to the unit vector x cannot increase more than twice:

$$\left\| x - \frac{p}{\|p\|} \right\| \leq \frac{2h}{k}.$$

Thus, the set

$$\mathcal{N}_{int} := \left\{ \frac{p}{\|p\|} \,:\, p = p(x) \text{ for some } x \in S_k \right\}$$

is a $2h/k$-net for S_k. How many different $p \in \mathbb{Z}^n$ we have to consider? Note that for any $x \in S_k$, the norm of $p(x)$ cannot be too large: since $\|x\| = 1$, $\operatorname{LCD}_{h,r_{3.36}}(x) < 2k$ and $4h/k < 2$, we get

$$\|p(x)\| \leq \operatorname{LCD}_{h,r_{3.36}}(x) + h < 3k.$$

Hence, all vectors $p \in \mathbb{Z}^n$ in the definition of $\mathcal{N}_{int}$ belong to the Euclidean ball of radius $3k$ centered at the origin. Standard volumetric argument shows that there are at most $(1 + Ck/\sqrt{n})^n$ integer points in this ball for a sufficiently large constant $C > 0$. Recall that $k \geq q_{3.36}\sqrt{n}/2$, whence

$$|\mathcal{N}_{int}| \leq \left(1 + \frac{Ck}{\sqrt{n}}\right)^n \leq \left(\frac{kL}{\sqrt{n}}\right)^n$$

for an appropriate number $L = L(\theta, \rho) > 0$. The net $\mathcal{N}_{int}$ does not have to be contained in S_k. But, by a standard argument, we can "replace" $\mathcal{N}_{int}$ with a $4h/k$-net of the same cardinality, and with elements from the set S_k. □

Together with Theorem 3.30, the above lemma gives

Lemma 3.38. *For any $\theta, \rho \in (0,1)$ there is $L_{3.38} = L_{3.38}(\theta, \rho) \geq 1$ such that for every $h \geq 1$ and $k > 0$ there is a finite subset $\mathcal{N} \subset S_k$ of cardinality at most $\left(kL_{3.38}/\sqrt{n}\right)^n$ with the following property. The event*

$$\left\{ \textit{For every } y \in S_k \textit{ there is } y' = y'(y) \in \mathcal{N} \textit{ such that } \|A(y - y')\| \leq hL_{3.38}\sqrt{n}/k \right\}$$

has probability at least $1 - 4\exp(-n/32)$.

Now, we can prove

Theorem 3.39. *Let α be a centered random variable of unit variance such that $\mathcal{L}(\alpha, \widetilde{v}) \leq \widetilde{u}$ for some $\widetilde{v} > 0$ and $\widetilde{u} \in (0,1)$. Then there exist $q, s, w, r > 0$ depending only on $\widetilde{v}, \widetilde{u}$ with the following property: let $X^1, X^2 \ldots, X^{n-1}$ be random n-dimensional vectors whose coordinates are jointly independent copies of α. Consider any random unit vector X^**

orthogonal to $\{X^1, X^2, \dots, X^{n-1}\}$*. Then*

$$\mathbb{P}\big\{\,\mathrm{LCD}_{s\sqrt{n},r}(X^*) < \exp(qn)\big\} \le 2\exp(-wn).$$

Proof. Without loss of generality, we can assume that n is a large number and that $\widetilde{v} \le 1$. Denote by A' the $n-1 \times n$ matrix with rows $X^1, X^2, \dots, X^{n-1}$. Then, by the definition of X^*, we have $A'X^* = 0$ almost surely. Let $\theta_{3.32}$ and $u_{3.32}$ be defined as in Remark 3.33 (with A' replacing A). Then, by Proposition 3.32 and Remark 3.33, we have

$$\mathbb{P}\big\{X^* \in \mathrm{Comp}_n(\theta_{3.32}, \theta_{3.32})\big\} \le 5{u_{3.32}}^n \le \exp(-wn)$$

for $w > 0$ such that, say, $\exp(-2w) > u_{3.32}$, and provided that n is large. Thus, it is enough to prove that

$$\mathbb{P}\big\{\,\mathrm{LCD}_{s\sqrt{n},r}(X^*) < \exp(qn), X^* \in \mathrm{Incomp}_n(\theta_{3.32}, \theta_{3.32})\big\} \le \exp(-wn)$$

for small enough r, w, s, q depending only on $\widetilde{v}, \widetilde{u}$. We start by defining $r := r_{3.36}(\theta_{3.32}, \theta_{3.32})$. Note that, by Lemma 3.36, we have

$$\mathrm{Incomp}_n(\theta_{3.32}, \theta_{3.32}) \subset \big\{x \in S^{n-1} :\ \mathrm{LCD}_{s\sqrt{n},r}(x) \ge q_{3.36}\sqrt{n}\big\}$$

for *any* $s > 0$, and, in particular for s defined by $s := \frac{\widetilde{v}r}{4L_{3.38}^2 L_{3.35}}$, where $L_{3.38} = L_{3.38}(\theta_{3.32}, \theta_{3.32})$ and $L_{3.35}$ are taken from Lemmas 3.38 and 3.35, respectively, and $q_{3.36} = q_{3.36}(\theta_{3.32}, \theta_{3.32})$. Let us emphasize that no vicious cycle is created here in regard to interdependence between s and r. Finally, we let $q := 2s^2(1-\widetilde{u})$ (w will be defined at the very end of the proof).

We will make use of representation (3.9) of the set $\mathrm{Incomp}_n(\theta_{3.32}, \theta_{3.32})$. Denote

$$\mathcal{K} := \big\{2^i :\ i \in [\log_2(q_{3.36}\sqrt{n}) - 1, qn/\ln 2] \cap \mathbb{N}\big\}.$$

Then, in view of Lemma 3.36, we have

$$\big\{x \in \mathrm{Incomp}_n(\theta_{3.32}, \theta_{3.32}) :\ \mathrm{LCD}_{s\sqrt{n},r}(x) < \exp(qn)\big\} \subset \bigsqcup_{k \in \mathcal{K}} S_k.$$

It is sufficient to prove that

$$\mathbb{P}\{X^* \in S_k\} \le 5\exp(-n/32) \quad \text{ for all } k \in \mathcal{K}. \tag{3.10}$$

Indeed, since $|\mathcal{K}| < qn$, the union bound over $\mathcal{K}$ will conclude the theorem.

In turn, (3.10) will follow as long as we show that

$$\mathbb{P}\{A'x = 0 \text{ for some } x \in S_k\} \le 5\exp(-n/32) \quad \text{for all } k \in \mathcal{K}.$$

Fix for a moment any $k \in \mathcal{K}$ and let $\mathcal{N}_k$ be the subset of S_k of cardinality at most $(kL_{3.38}/\sqrt{n})^n$, constructed in Lemma 3.38 (with $h := s\sqrt{n}$). Further, take $\varepsilon := \frac{\widetilde{v}r\sqrt{n}}{2kL_{3.38}L_{3.35}}$. Note that, in view of the definition of q and $\mathcal{K}$, we have $k \le \exp(2s^2(1-\widetilde{u})n)$. Hence, for n large enough, ε satisfies the condition of Lemma 3.35:

$$\varepsilon \ge \widetilde{v} \cdot \max\Big(\frac{1}{k}, \exp\big(-2s^2(1-\widetilde{u})n\big)\Big) \ge \widetilde{v} \cdot \max\Big(\frac{1}{\mathrm{LCD}_{h,r}(x)}, \exp(-2(1-\widetilde{u})h^2)\Big).$$

Hence,

$$\begin{aligned}
\mathbb{P}\big\{\|A'y\| \ge \varepsilon\sqrt{n} \text{ for all } y \in \mathcal{N}_k\big\} &\ge 1 - |\mathcal{N}_k|(L_{3.35}\varepsilon/r)^{n-1} \\
&\ge 1 - \Big(\frac{kL_{3.38}}{\sqrt{n}}\Big)^n \Big(\frac{L_{3.35}\varepsilon}{r}\Big)^{n-1} \\
&\ge 1 - \frac{kL_{3.38}}{\sqrt{n}} \cdot \Big(\frac{\widetilde{v}}{2}\Big)^{n-1} \\
&\ge 1 - 2^{-n}\exp(2s^2(1-\widetilde{u})n),
\end{aligned}$$

where the last relation follows by the assumption $\widetilde{v} \le 1$. Finally, note that, since $s \le 1/4$, the last quantity is bounded from below by $1 - 2^{-n/2}$. Applying the definition of $\mathcal{N}_k$ in Lemma 3.38 and noticing that $hL_{3.38}\sqrt{n}/k \le \varepsilon\sqrt{n}/2$, we get

$$\mathbb{P}\big\{\|A'y\| \ge \varepsilon\sqrt{n}/2 \text{ for all } y \in S_k\big\} \ge 1 - 4\exp(-n/32) - 2^{-n/2} \ge 1 - 5\exp(-n/32).$$

This proves (3.10) and implies the result. □

Proof of Theorem 3.1. Without loss of generality, the dimension n is large. Let $A = (a_{ij})$ be an $n \times n$ random matrix with i.i.d. centered entries with unit variance such that for some $\widetilde{v} > 0$ and $\widetilde{u} \in (0,1)$ we have $\mathcal{L}(a_{ij}, \widetilde{v}) \le \widetilde{u}$. We define $\theta := \theta_{3.32}(\widetilde{v}, \widetilde{u})$ and $v := v_{3.32}(\widetilde{v}, \widetilde{u})$, where $\theta_{3.32}, v_{3.32}$ are taken from Proposition 3.32, and let q, s, w, r be as in Theorem 3.39 (with respect to $\widetilde{v}, \widetilde{u}$). We will prove a small ball probability bound for $s_n(A)$.

It is sufficient to consider the parameter domain $\varepsilon \in \big(\theta\widetilde{v}\exp(-qn), 1\big]$. We have

$$\begin{aligned}
\mathbb{P}\big\{s_n(A) < \varepsilon n^{-1/2}\big\} &\le \mathbb{P}\Big\{\inf_{y \in \mathrm{Comp}_n(\theta,\theta)} \|Ay\| < v\sqrt{n}\Big\} + \mathbb{P}\Big\{\inf_{y \in \mathrm{Incomp}_n(\theta,\theta)} \|Ay\| < \varepsilon n^{-1/2}\Big\} \\
&\le 5{u_{3.32}}^n + \mathbb{P}\Big\{\inf_{y \in \mathrm{Incomp}_n(\theta,\theta)} \|Ay\| < \varepsilon n^{-1/2}\Big\},
\end{aligned}$$

where we have applied Proposition 3.32. Further, by Proposition 3.22, we have

$$\mathbb{P}\Big\{\inf_{y\in\mathrm{Incomp}_n(\theta,\theta)}\|Ay\| < \varepsilon n^{-1/2}\Big\} \le \frac{1}{\theta}\mathbb{P}\Big\{\Big|\sum_{i=1}^{n} X_i^* a_{in}\Big| < \frac{\varepsilon}{\theta}\Big\},$$

where X^* denotes a random unit normal vector to the span of the first $n-1$ columns of A. In view of Theorem 3.29, this last relation implies

$$\begin{aligned}\mathbb{P}\Big\{\inf_{y\in\mathrm{Incomp}_n(\theta,\theta)}\|Ay\| < \varepsilon n^{-1/2}\Big\} &\le \theta^{-1}\mathbb{P}\big\{\mathrm{LCD}_{s\sqrt{n},r}(X^*) < \theta\widetilde{v}\varepsilon^{-1}\big\}\\ &\quad + \frac{C_{3.29}\varepsilon}{\theta\widetilde{v}r\sqrt{1-\widetilde{u}}} + C_{3.29}\exp\big(-2s^2(1-\widetilde{u})n\big).\end{aligned}$$

Finally, noticing that $\theta\widetilde{v}\varepsilon^{-1} \le \exp(qn)$ and applying Theorem 3.39, we get

$$\begin{aligned}\mathbb{P}\Big\{\inf_{y\in\mathrm{Incomp}_n(\theta,\theta)}\|Ay\| < \varepsilon n^{-1/2}\Big\} &\le 2\theta^{-1}\exp(-wn)\\ &\quad + \frac{C_{3.29}\varepsilon}{\theta\widetilde{v}r\sqrt{1-\widetilde{u}}}\\ &\quad + C_{3.29}\exp\big(-2s^2(1-\widetilde{u})n\big).\end{aligned}$$

Together with an estimate for the compressible vectors, this implies the result. □

CHAPTER 4

Operator norm: regularization

4.1 Motivation and main results

When a certain mathematical or scientific structure fails to meet reasonable expectations, one often wonders: is this a *local or global problem*? In other words, is the failure caused by some small, localized part of the structure, and if so, can this part be identified and repaired? Or, alternatively, is the structure entirely, globally bad? Many results in mathematics can be understood as either local or global statements. For example, not every measurable function $f : \mathbb{R} \to \mathbb{R}$ is continuous, but Lusin's theorem implies that f can always be made continuous by changing its values on a set of arbitrarily small measure. Thus, imposing continuity is a local problem. On the other hand, a continuous function may not be differentiable, and there even exist continuous and nowhere differentiable functions. Thus imposing differentiability may be a global problem. In statistics, the notion of *outliers* – small, pathological subsets of data, the removal of which makes data better – points to local problems.

So, is bounding the norm of a random matrix a local or a global problem? To be specific, consider $n \times n$ random matrices A with independent and identically distributed (i.i.d.) entries. Recall that the *operator norm* of A is defined by considering A as a linear operator on $\mathbb{R}^n$ equipped with the Euclidean norm $\|\cdot\|_2$, i.e.

$$\|A\| = \max_{x \neq 0} \frac{\|Ax\|_2}{\|x\|_2}.$$

Suppose that the entries of A have zero mean and bounded fourth moment, i.e. $\mathbb{E}A_{ij}^4 = O(1)$. Then, as it was shown in [YBK88],

$$\|A\| = (2 + o(1))\sqrt{n}$$

with high probability. Note that the order $\sqrt{n}$ is the best we can generally hope for. Indeed, if the entries of A have unit variance, then the typical magnitude of the Euclidean norm of a row of A is $\sim \sqrt{n}$, and the operator norm of A can not be smaller than that. Moreover, by [BSY88, Sil89] the bounded fourth moment assumption is nearly necessary for the bound

$$\|A\| = O(\sqrt{n}). \tag{4.1}$$

To put this precisely, consider an infinite array of i.i.d. random variables $\{A_{ij} : i, j = 1, 2, \dots\}$, and the sequence of $n \times n$ random matrices $A_n := (A_{ij})_{i,j=1,\dots,n}$. Then for almost surely convergence fourth moment is necessary and sufficient ($\mathbb{E}A_{11}^4 = \infty$ implies a.s. $\limsup \|A_n\|^2/n = \infty$, see [BSY88]). For convergence in probability, zero mean $\mathbb{E}A_{11} = 0$ and the weak fourth moment ($n^2\mathbb{P}\{|A_{11}| \geq n\} = o(1)$) are necessary and sufficient, see [Sil89].

A number of quantitative and more general versions of these bounds are known [Seg00, Lat05, Vu05, BVH16, vH17a, vH17b].

Now let us postulate nothing at all about the distribution of the i.i.d. entries of A. It still makes sense to ask: *is enforcing the ideal bound* (4.1) *for random matrices a local or a global problem*? That is, can we enforce the bound (4.1) by modifying the entries in a small submatrix of A? We have shown that this is possible if and only if the entries of A have zero moment and finite variance. The "if" part is covered by the following theorem.

Theorem 4.1 (Local problem)**.** *Consider an $n \times n$ random matrix A with i.i.d. entries that have zero mean and unit variance, and let $\varepsilon \in (0, 1/6]$. Then, with probability at least $1 - 7\exp(-\varepsilon n/12)$, there exists an $\varepsilon n \times \varepsilon n$ submatrix of A such that replacing all of its entries with zero leads to a well-bounded matrix $\tilde{A}$:*

$$\|\tilde{A}\| \leq \frac{C \ln \varepsilon^{-1}}{\sqrt{\varepsilon}} \cdot \sqrt{n},$$

where C is a sufficiently large absolute constant.

Remark 4.2 (Optimality). The dependence on ε in Theorem 4.1 is best possible up to the $\ln \varepsilon^{-1}$ factor. To see this, let $p := 2\varepsilon/n$ and suppose A_{ij} take values $\pm 1/\sqrt{p}$ with probability $p/2$ each and value 0 with probability $1-p$. Then A_{ij} have zero mean and unit variance as required. The expected number of non-zero entries in A equals $pn^2 = 2\varepsilon n$. Thus the number of the rows of A containing these entries is bigger than εn with high probability. (This is a standard observation about the balls-into-bins model.) Therefore, no $\varepsilon n \times \varepsilon n$ submatrix can contain all the non-zero entries of A. In other words, $\tilde{A}$ must contain at least one non-zero entry of A, and thus it has magnitude

$$\|\tilde{A}\| \geq \frac{1}{\sqrt{p}} \gtrsim \frac{\sqrt{n}}{\sqrt{\varepsilon}}.$$

This shows that the dependence on ε in Theorem 4.1 is almost optimal.

By rescaling, a more general version of Theorem 4.1 holds for any finite variance of the entries. The two main assumptions in this theorem – mean zero and finite variance – are necessary in Theorem 4.1. Without either of them, the problem becomes global in a strong sense: the desired $O(\sqrt{n})$ bound can not be achieved even after modifying a *large* submatrix. This is the content of the following result.

Theorem 4.3 (Global problem)**.** *Consider an $n \times n$ random matrix A_n whose entries are i.i.d. copies of a random variable that has either nonzero mean or infinite second moment,*[1] *and let $\varepsilon \in (0,1)$. Then*

$$\min \frac{\|\tilde{A}_n\|}{\sqrt{n}} \to \infty \quad \textit{as} \quad n \to \infty$$

almost surely. Here the minimum is with respect to the matrices $\tilde{A}_n$ obtained by any modification of any $\varepsilon n \times \varepsilon n$ submatrix of A_n.

It should be noted that while Theorem 4.1 becomes harder for smaller ε, Theorem 4.3 becomes harder for larger ε, those near 1.

We prove Theorem 4.3 in Section 4.3. The argument is considerably simpler than for Theorem 4.1. Indeed, the nonzero mean forces the sum of the entries of $\tilde{A}_n$ to be $\gtrsim n^2$, and the infinite second moment forces the Frobenius norm of $\tilde{A}_n$ (the square root of the sum of the entries squared) to be $\gg n^2$ with high probability. Either of these two bounds can be easily used to show that the operator norm of $\tilde{A}_n$ is $\gg \sqrt{n}$.

4.1.1 Overview of the proof

The approach to the proof of Theorem 4.1 utilizes and advances the ideas described in the previous chapter (wile proving invertibility Theorem 3.1), and combines them with the methods developed recently in [LLV17]. We first control the cut norm of A and then pass to the operator norm using Grothendieck-Pietsch factorization. Let us describe these steps in more detail.

[1] Although this is a minor terminological distinction, in this theorem we prefer to talk about second moment rather than variance. This is because the second moment $\mathbb{E}X^2$ of a random variable X is always defined in the extended real line, while the variance $\mathrm{Var}(X) = \mathbb{E}(X - \mathbb{E}X)^2$ is undefined if the mean $\mathbb{E}X$ is infinite.

The operator norm of a matrix A, as we already mentioned, is defined by considering A as a linear operator on the (finite dimensional) space ℓ_2, i.e.

$$\|A\| = \|A : \ell_2 \to \ell_2\|.$$

Rather than bounding the operator norm of a random matrix A directly, we shall compare it with two simpler norms,

$$\|A\|_{\infty\to 2} = \|A : \ell_\infty \to \ell_2\| = \max_{x\neq 0} \frac{\|Ax\|_2}{\|x\|_\infty}$$

and

$$\|A\|_{2\to\infty} = \|A : \ell_2 \to \ell_\infty\| = \max_{x\neq 0} \frac{\|Ax\|_\infty}{\|x\|_2}.$$

The simplest of the three is the $2 \to \infty$ norm. A quick check reveals that it equals the maximum Euclidean norm of the rows A_i^T of A:

$$\|A\|_{2\to\infty} = \max_{i\in[n]} \|A_i\|_2. \tag{4.2}$$

The next simplest norm is $\infty \to 2$, which can be conveniently computed as

$$\|A\|_{\infty\to 2} = \max_{x\in\{-1,1\}^n} \|Ax\|_2. \tag{4.3}$$

This norm is equivalent within a constant factor to the *cut norm* from the computer science literature [BJR10, AN04], where the maximum is taken over $\{0,1\}^n$. The hardest of the three is the operator norm,

$$\|A\| = \max_{x\in S^{n-1}} \|Ax\|_2. \tag{4.4}$$

To see why the difficulty in bounding these norms rises this way, note that one has to control n random variables in (4.2), 2^n random variables in (4.3), and infinitely many random variables in (4.4).

How large do we expect the three norms to be for random matrices? For a simple example, let us first consider a Gaussian random matrix A with i.i.d. $N(0,1)$ entries. Then it is not difficult to check that

$$\|A\|_{2\to\infty} \sim \sqrt{n}, \quad \|A\|_{\infty\to 2} \sim n, \quad \|A\| \sim \sqrt{n}. \tag{4.5}$$

Indeed, note that the rows of A have Euclidean norms $\sqrt{n}$ on average, so the bound on the

$2 \to \infty$ norm follows by union bound and using Gaussian concentration. The bound on the $\infty \to 2$ norm follows from (4.3) by using Gaussian concentration for the normal random vector Ax and taking the union bound over $\{-1, 1\}^n$. The bound on the operator norm is a non-asymptotic version of Bai-Yin's law, see e.g. [Ver12, Theorem 5.32].

One might wonder if (4.5) holds not only in the Gaussian case but generally for random matrices A with i.i.d. entries that have zero mean and unit variance. In particular, it would be wonderful if the three norms were always related to each other as follows:

$$\|A\| \lesssim \frac{\|A\|_{\infty\to 2}}{\sqrt{n}} \lesssim \|A\|_{2\to\infty} \lesssim \sqrt{n}. \tag{4.6}$$

This, however, would be too optimistic to expect, since the bound $\|A\| \lesssim \sqrt{n}$ cannot hold without higher moments assumptions as we mentioned in Section 4.1. Nevertheless, we will obtain a version of (4.6) after removal a small fraction of rows of A. With high probability, we will be able to find subsets of rows $J_1 \subset J_2 \subset J_3$ with cardinalities $|J_i| \leq \varepsilon n$ and such that

$$\|A_{J_3^c}\| \lesssim \frac{\|A_{J_2^c}\|_{\infty\to 2}}{\sqrt{n}} \lesssim \|A_{J_1^c}\|_{2\to\infty} \lesssim \sqrt{n}. \tag{4.7}$$

where the inequalities hide a factor that depends on ε.

The first step in proving (4.7) is to find a small set J_1 with $|J_1| \lesssim \varepsilon n$ and such that

$$\|A_{J_1^c}\|_{2\to\infty} \lesssim \sqrt{n} \tag{4.8}$$

with high probability. In other words, we would like to bound all rows of A simultaneously by $O(\sqrt{n})$ after removing a few columns of A. To show this we first focus on one row, where we need to bound a sum of independent random variables (the squares of the row's entries). In Theorem 4.6 we show how to bound sums of independent random variables almost surely by gently *damping* the summands. Damping, or reweighting down, is a softer operation than removing entries. It allows us to treat in Section 4.2.2 all columns simultaneously without much effort, thus proving (4.8). The argument in this step is similar to the one developed to fit a random vector into an ℓ_p^n-ball (Section 3.2.2). We still need to improve the dependence between the number of removed columns and the resulting $2 \to \infty$ norm; this will ultimately lead to the optimal dependence on ε in Theorem 4.1.

At the next step, we extend J_1 to a bigger set of rows J_2 with $|J_2| \lesssim \varepsilon n$ and so that

$$\|A_{J_2^c}\|_{\infty\to 2} \lesssim n. \tag{4.9}$$

Suppose for a moment that we are not concerned about removal of any columns. It is not too hard to show the general bound

$$\mathbb{E}\|A\|_{\infty\to 2} \lesssim \sqrt{n}\mathbb{E}\|A\|_{2\to\infty}, \tag{4.10}$$

for a random matrix A with independent, mean zero entries; we prove this in Lemma 4.8. However, this bound is not very helpful in our situation. We need to work with the matrix $A_{J_1^c}$ instead of A, which is not trivial: the removal of the columns in J_1 that we did in the first step made the entries of $A_{J_1^c}$ dependent. In Lemma 4.9, we first prove a variant of (4.10) for $A_{J_1^c}$ under an additional symmetry assumption on the distribution of the entries of A. Then we manage to remove this assumption with a delicate symmetrization argument, which we develop in the rest of Section 4.2.2, with the final result being Theorem 4.13. The general idea of this step again follows the ideas from the previous chapter (proof of Theorem 3.2). However here we need to produce much more delicate symmetrization argument to obtain (4.9) with a *logarithmic* dependence on ε.

Next, we pass from $\infty \to 2$ norm to the operator norm in Section 4.2.2. This is done by using Grothendieck-Pietsch factorization (Theorem 4.14), a result that yields the first inequality in (4.7) for completely arbitrary, even non-random, matrices. This reasoning was recently used in a similar context in [LLV17].

The argument we just described works under the additional assumption that the entries of A be $O(\sqrt{n})$ almost surely. To be specific, such boundedness assumption is needed to make the damping argument in Step 1 work with mild, logarithmic dependence on ε. The contribution of the entries that are larger than $\sqrt{n}$ are controlled in Section 4.2.4 by showing that there can not be too many of them. The unit variance assumption implies that there are $O(1)$ such large entries per column on average. This does not mean, of course, that all columns will have $O(1)$ large entries with high probability; in fact there could be columns with $\sim \log n/\log\log n$ large entries. But we will check in Lemma 4.17 that the number of such heavy columns is small; removing them will lead to the desired bound $O(\sqrt{n})$ on the operator norm for the matrix with large entries. We develop this argument in Proposition 4.20 and Corollary 4.22, and derive the full strength of Theorem 4.1 in the end of Section 4.2.4.

Theorem 4.3 is proved in Section 4.3.

4.2 Local problem

4.2.1 Damping a sum of independent random variables

Remark 4.4. Both results and methods of this section follow closely the ones of Section 3.2.2. The only important improvement is that we obtain better (logarithmic) dependence on ε in (4.13) in trade of the additional boundedness assumption. For the sake of integrity of the exposition, we decided to keep this section self-contained, and re-introduce everything needed for the further sections with the most convenient notations.

Let $X_1, \ldots, X_n$ be non-negative i.i.d. random variables with $\mathbb{E}X_i \leq 1$. The linearity of expectation gives the trivial bound

$$\mathbb{E}\sum_{i=1}^{n} X_i \leq n.$$

Here we will be interested in a stronger result – that the sum be $O(n)$ *almost surely* instead of in expectation. To do this, we will be looking for random weights

$$W_1, \ldots, W_n \in [0,1]$$

that make the "damped" sum satisfy

$$\sum_{j=1}^{n} W_j X_j = O(n) \quad \text{almost surely.}$$

To make the damping as gentle as possible, we are looking for largest possible weights W_i, hopefully very close to 1.

To get started, let us consider the simple case where $n = 1$ and try to damp one random variable.

Lemma 4.5 (Damping a random variable)**.** *Let X be a random variable such that*

$$X \geq 0 \quad \textit{and} \quad \mathbb{E}X \leq 1.$$

Let $\varepsilon \in (0,1)$. There exists a random variable W taking values in $[0,1]$ and such that

$$XW \leq \varepsilon^{-1} \quad \textit{almost surely;} \tag{4.11}$$

$$1 \leq \mathbb{E}W^{-1} \leq 1+\varepsilon. \tag{4.12}$$

Proof. Fix a level $L \geq 1$ whose value we will choose later, and define

$$W := \min(1, L/X).$$

To check (4.11), we have

$$XW = \min(X, L) \leq L \quad \text{almost surely}.$$

Next, the lower bound in (4.12) holds trivially since $W \leq 1$. For the upper bound, we have

$$\mathbb{E}W^{-1} = \mathbb{E}\max(1, X/L) \leq \mathbb{E}(1 + X/L) \leq 1 + \frac{1}{L},$$

where we used the assumption that $\mathbb{E}X \leq 1$. Setting $L = \varepsilon^{-1}$ completes the proof. $\square$

Now let us address the damping problem for general number n of random variables, which we described in the beginning of this section. Applying Lemma 4.5 for each random variable X_i, we get weights W_i such that

$$\sum_{j=1}^{n} W_j X_j \leq \varepsilon^{-1} n \quad \text{almost surely;}$$
$$1 \leq \mathbb{E}\Big(\prod_{j=1}^{n} W_j\Big)^{-1} \leq (1+\varepsilon)^n = 1 + O(\varepsilon n)$$

for small ε. Note that this would be exactly the result proved in Proposition 3.4 in the case $p = 2$. We will now considerably improve both these bounds, making only one mild extra assumption that $X_i = O(n)$ almost surely.

Theorem 4.6 (Damping a sum of random variables)**.** *Let $X_1, \ldots, X_n$ be i.i.d. random variables such that*

$$0 \leq X_j \leq Kn \quad \textit{and} \quad \mathbb{E}X_j \leq 1$$

for some $K \geq 1$. Let $\varepsilon \in (0, 1/2)$. There exist random variables $W_1, \ldots, W_n$ taking values in $[0, 1]$ and such that

$$\sum_{j=1}^{n} W_j X_j \leq CK\log(\varepsilon^{-1}) \cdot n \quad \textit{almost surely;} \tag{4.13}$$

$$1 \leq \mathbb{E}\Big(\prod_{j=1}^{n} W_j\Big)^{-1} \leq 1 + \varepsilon. \tag{4.14}$$

Proof. **Step 1: Bernoulli distribution.** Let us first prove the theorem in the particular case where X_j are scaled Bernoulli random variables. Assume that X_j can take values q and 0, and

$$\mathbb{P}\left\{X_j = q\right\} = p \geq \frac{1}{Kn}. \tag{4.15}$$

Let ν denote the (random) number of nonzero X_j's:

$$\nu := \left|\{j : \ X_j \neq 0\}\right|, \quad \text{then} \quad \mathbb{E}\nu = pn.$$

Here is how we will define the weights W_j. If $X_j = 0$ then clearly there is no need to damp X_j so put $W_j = 1$. The same applies if the number ν of non-zero X_j's does not significantly exceed its expectation pn. Otherwise we damp all terms by the same amount $W_j \sim pn/\nu$. Formally, we fix some parameter $L = L(K, \varepsilon)$ whose value we will determine later, and set

$$W_j := \begin{cases} 1, & \text{if } \nu \leq Lpn \text{ or } X_j = 0 \\ Lpn/\nu, & \text{if } \nu > Lpn \text{ and } X_j \neq 0. \end{cases}$$

Let us check (4.13). In the event when $\nu \leq Lpn$, we have

$$\sum_{j=1}^{n} W_j X_j = \sum_{j=1}^{\nu} 1 \cdot q = q\nu \leq qLpn = Ln \cdot \mathbb{E}X_1.$$

And in the event when $\nu > Lpn$, we have

$$\sum_{j=1}^{n} W_j X_j = \sum_{j=1}^{\nu} \frac{Lpn}{\nu} \cdot q = Lpnq = Ln \cdot \mathbb{E}X_1.$$

as before. Thus, we showed that

$$\sum_{j=1}^{n} W_j X_j \leq Ln \cdot \mathbb{E}X_1 \leq Ln \quad \text{almost surely.} \tag{4.16}$$

Let us now check (4.14). Since the lower bound is trivial, we will only have to check the upper bound. We will again split the calculation into two cases based on the size of ν. If $\nu \leq Lpn$ then all $W_j = 1$, so we trivially get

$$E_- := \mathbb{E}\Big(\prod_{j=1}^{n} W_j\Big)^{-1} \mathbb{1}_{\{\nu \leq Lpn\}} \leq 1.$$

If $\nu > Lpn$, then the definition of W_j gives

$$E_+ := \mathbb{E}\Big(\prod_{j=1}^n W_j\Big)^{-1} \mathbb{1}_{\{\nu > Lpn\}} = \mathbb{E}\Big(\frac{\nu}{Lpn}\Big)^{\nu} \mathbb{1}_{\{\nu > Lpn\}}$$

$$= \sum_{k=\lceil Lpn \rceil+1}^{n} \Big(\frac{k}{Lpn}\Big)^k \mathbb{P}\left\{\nu = k\right\}.$$

Since $\nu \sim \mathrm{Binom}(n, p)$, we have

$$\mathbb{P}\left\{\nu = k\right\} = \binom{n}{k} p^k \leq \Big(\frac{enp}{k}\Big)^k,$$

using a standard consequence of Stirling's approximation. Thus

$$E_+ \leq \sum_{k=\lceil Lpn \rceil+1}^{n} \Big(\frac{e}{L}\Big)^k \leq \Big(\frac{e}{L}\Big)^{Lpn},$$

provided that $L \geq 10$. Thus we showed that

$$\mathbb{E}\Big(\prod_{j=1}^n W_j\Big)^{-1} \leq E_- + E_+ \leq 1 + \Big(\frac{e}{L}\Big)^{Lpn} \leq 1 + \Big(\frac{e}{L}\Big)^{L/K} \tag{4.17}$$

where in the last step we used the assumption that $p \geq 1/Kn$ that we made in (4.15).

Now that we have the bounds (4.16) and (4.17), it is enough to choose

$$L := CK \log\Big(\frac{1}{\varepsilon}\Big)$$

which implies that $E \leq 1 + \varepsilon$. The proof for the Bernoulli distribution is complete.

Step 2. General distribution. Let us now now prove the theorem in full generality. First we discretize the distribution of X_j using Lemma 2.15. This result requires X_j to be continuous, which can be arranged by a standard approximation argument. For example, we can add a small Gaussian independent component to X_j and then let the variance of this component go to zero. Taking into account Remarks 2.16 and 2.17, we obtain independent, non-negative random variables X_j' that satisfy $\mathbb{E}X_j' \leq 4$ and such that

$$X_j \leq X_j' = \sum_{k=1}^{\kappa} X_{jk}.$$

Here X_{jk} are independent random variables; each X_{jk} can take values q_k and 0, and

$$\mathbb{P}\left\{X_{jk} = q_k\right\} = p_k$$

with

$$p_k = 2^{-k} \geq \frac{1}{Kn} \text{ for } k < \kappa, \quad p_\kappa = \frac{1}{Kn}. \tag{4.18}$$

The argument will be similar to step 1 of the proof. For each level k we let ν_k denote number of non-zero X_{jk}'s:

$$\nu_k := \left|\{j :\ X_{jk} \neq 0\}\right|, \quad \text{then} \quad \mathbb{E}\nu = p_k n.$$

Again, for each level k define the weights W_{jk} like in step 1:

$$W_{jk} := \begin{cases} 1, & \text{if } \nu_k \leq Lp_k n \text{ or } X_{jk} = 0 \\ Lp_k n/\nu_k, & \text{if } \nu_k > Lp_k n \text{ and } X_{jk} \neq 0. \end{cases}$$

Then we set

$$W_j := \prod_{k=1}^{\kappa} W_{jk}, \quad j = 1, \ldots, n.$$

Let us check (4.13). We have

$$\sum_{j=1}^{n} W_j X_j \leq \sum_{j=1}^{n} W_j X_j' = \sum_{j=1}^{n} \sum_{k=1}^{\kappa} W_j X_{jk} \leq \sum_{k=1}^{\kappa} \sum_{j=1}^{n} W_{jk} X_{jk}, \tag{4.19}$$

since $W_j \leq W_{jk}$ by construction. Now, for each level k, we can use step 1 of the proof, where we showed in (4.16) that

$$\sum_{j=1}^{n} W_{jk} X_{jk} \leq Ln \cdot \mathbb{E}X_{1k}.$$

Substituting into (4.19), we obtain

$$\sum_{j=1}^{n} W_j X_j \leq Ln \cdot \sum_{k=1}^{\kappa} \mathbb{E}X_{1k} = Ln \cdot \mathbb{E}X_1' \leq 5Ln \tag{4.20}$$

by construction.

Let us now check (4.14). The lower bound is trivial, and we will only have to check the upper bound. For each level k, we can use step 1 of the proof, where we showed in (4.17)

that
$$\mathbb{E}\Big(\prod_{j=1}^{n} W_{jk}\Big)^{-1} \le 1 + \Big(\frac{e}{L}\Big)^{Lp_k n} \le 1 + e^{-Lp_k n},$$

which is true as long as $L \ge 10$. Then, by construction we have

$$\begin{aligned}
\mathbb{E}\Big(\prod_{j=1}^{n} W_j\Big)^{-1} &= \mathbb{E}\prod_{k=1}^{\kappa}\Big(\prod_{j=1}^{n} W_{jk}\Big)^{-1} \\
&= \prod_{k=1}^{\kappa}\mathbb{E}\Big(\prod_{j=1}^{n} W_{jk}\Big)^{-1} \quad \text{(by independence)} \\
&\le \prod_{k=1}^{\kappa}\big(1 + e^{-Lp_k n}\big) \le \exp\Big(\sum_{k=1}^{\kappa} e^{-Lp_k n}\Big)
\end{aligned}$$

where in the last step we used the inequality $1 + x \le e^x$. Recall from (4.18) that the exponents p_k form a decreasing geometric progression with values 2^{-k} until the last (smallest) term of order $1/Kn$. So this last term dominates the sum $\sum_{k=1}^{\kappa} e^{-Lp_k n}$, and we obtain

$$\mathbb{E}\Big(\prod_{j=1}^{n} W_j\Big)^{-1} \le \exp(2e^{-L/2K}). \tag{4.21}$$

Now that we have the bounds (4.20) and (4.21), it is enough to choose

$$A_{ij}L := C_{4.6} K \log\Big(\frac{1}{\varepsilon}\Big)$$

with $C_{4.6} \ge 6K$ and the right hand side of (4.21) will be bounded by

$$\exp(2\varepsilon^3) \le \exp(\varepsilon/2) \le 1 + \varepsilon,$$

as claimed. The proof of the theorem is complete. $\square$

4.2.2 Controlling the bounded entries: three matrix norms

In this section we prove Theorem 4.1 under the additional assumption that all entries A_{ij} of A are not too large. Specifically, let us assume that

$$|A_{ij}| \le \frac{\sqrt{n}}{2} \quad \text{almost surely.} \tag{4.22}$$

The $2 \to \infty$ norm of random matrices

Lemma 4.7 (Bounding $2 \to \infty$ norm by removing a few columns)**.** *Consider an $n \times n$ random matrix A with i.i.d. entries A_{ij} which have mean zero and at most unit variance and satisfy* (4.22)*. Let $\varepsilon \in (0, 1/2]$. Then with probability at least $1 - \exp(-\varepsilon n)$, there exists a subset $J \in [n]$ with cardinality $|J| \le \varepsilon n$ such that*

$$\|A_{J^c}\|_{2\to\infty} \le C\sqrt{\ln \varepsilon^{-1}} \cdot \sqrt{n}.$$

Proof. We apply Theorem 4.6 for the squares of the elements in each row of A, i.e. for the random variables $(a_{i1}^2, \dots, a_{in}^2)$. This gives us random weights $W_{ij} \in [0, 1]$ which satisfy for each $i \in [n]$ that

$$\sum_{j=1}^n W_{ij} A_{ij}^2 \le C \log(\varepsilon^{-1}) n \quad \text{a.s.}; \qquad \mathbb{E}\Big(\prod_{j=1}^n W_{ij}\Big)^{-1} \le \exp(\varepsilon).$$

To make the same system of weights work for all rows, we define

$$V_j := \prod_{i=1}^n W_{ij} \in [0, 1], \quad j \in [n].$$

Then obviously $V_j \le W_{ij}$ for every i, and so

$$\sum_{j=1}^n V_j A_{ij}^2 \le C \log(\varepsilon^{-1}) n \;\; \forall i \quad \text{a.s.}; \qquad \mathbb{E}\Big(\prod_{j=1}^n V_j\Big)^{-1} \le \exp(\varepsilon n). \tag{4.23}$$

We will remove from A the columns whose weights V_j are too small, namely those in

$$J := \{j \in [n] : \; V_j < e^{-2}\}.$$

Let us first check that

$$|J| \le \varepsilon n \quad \text{with probability at least } 1 - \exp(-\varepsilon n), \tag{4.24}$$

as we claimed in the lemma. Indeed, if $|J| > \varepsilon n$ then using that all $V_j \in [0, 1]$ we have

$$Z := \prod_{j=1}^n V_j \le \prod_{j \in J} V_j < e^{-2\varepsilon n}.$$

But the probability of this event can be bounded by Markov's inequality:

$$\mathbb{P}\left\{Z < e^{-2\varepsilon n}\right\} = \mathbb{P}\left\{Z^{-1} > e^{2\varepsilon n}\right\} \leq e^{-2\varepsilon n}\mathbb{E}Z^{-1} \leq e^{-\varepsilon n},$$

where in the last bound we used (4.23). This proves (4.24).

It remains to check that all rows B_i of the matrix $B = A_{[n]\times J_0^c}$ are bounded as claimed. We have

$$\begin{aligned}
\|B_i\|_2^2 = \sum_{j\in J^c} A_{ij}^2 &\leq e^2 \sum_{j\in J^c} V_j A_{ij}^2 \quad \text{(by definition of } J\text{)} \\
&\leq e^2 \sum_{j=1}^{n} V_j A_{ij}^2 \quad \text{(since all } V_j \leq 1\text{)} \\
&\leq e^2 C \ln(\varepsilon^{-1}) n \quad \text{(by (4.23))}.
\end{aligned}$$

Taking the square root of both sides completes the proof. □

From $2 \to \infty$ norm to $\infty \to 2$ norm

Now we will control the $\infty \to 2$ norm of a random matrix. Our first task is to bound the $\infty \to 2$ norm by the simpler $2 \to \infty$ norm. The resulting comparison inequalities are interesting in their own right; we state them in Lemmas 4.8 and 4.10. The ultimate result of this section is Theorem 4.13, which gives an optimal bound $O(n)$ on the $\infty \to 2$ norm of a random matrix after removing a small fraction of columns.

The first method is based on flipping the signs of the entries independently at random. Here is the main result of this section.

Lemma 4.8 (From $2 \to \infty$ to $\infty \to 2$)**.** *Let A be an $n \times n$ random matrix whose entries are independent, mean zero random variables. Then*

$$\mathbb{E}\|A\|_{\infty\to 2} \leq C\sqrt{n} \cdot \mathbb{E}\|A\|_{2\to\infty}.$$

Proof. Let ε_{ij} be independent Rademacher random variables (which are also independent of A) and consider the random matrix

$$\tilde{A} := (\varepsilon_{ij} A_{ij}).$$

A basic symmetrization inequality (see [LT13, Lemma 6.3]) yields

$$\mathbb{E}\|A\|_{\infty\to 2} \le 2\mathbb{E}\|\tilde{A}\|_{\infty\to 2}.$$

Condition on A; the randomness now rests in the random signs (ε_{ij}) only. It suffices to show that the conditional expectation satisfies

$$\mathbb{E}\|\tilde{A}\|_{\infty\to 2} \lesssim \sqrt{n}\cdot\|A\|_{2\to\infty}. \tag{4.25}$$

Recalling (4.3), we have

$$\|\tilde{A}\|_{\infty\to 2} = \max_{x\in\{-1,1\}^n}\|\tilde{A}x\|_2. \tag{4.26}$$

According to the matrix-vector multiplication, we can express $\|\tilde{A}x\|^2$ as a sum of independent random variables

$$\|\tilde{A}x\|_2^2 = \sum_{i=1}^n \xi_i^2 \quad \text{where} \quad \xi_i := \langle \tilde{A}_i, x\rangle = \sum_{j=1}^n \varepsilon_{ij}A_{ij}x_j.$$

Fix $x \in \{-1,1\}^n$. Using independence and (4.2), we get

$$\mathbb{E}\xi_i^2 = \sum_{j=1}^n (A_{ij}x_{ij})^2 = \sum_{j=1}^n A_{ij}^2 \le \|A\|_{2\to\infty}^2,$$

so

$$\mathbb{E}\sum_{i=1}^n \xi_i^2 \le n\|A\|_{2\to\infty}^2. \tag{4.27}$$

Moreover, the standard concentration results ([Ver12, Lemma 5.9]) show that each ξ_i is a subgaussian random variable, and we have

$$\|\xi_i\|_{\psi_2}^2 = \Big\|\sum_{j=1}^n \varepsilon_{ij}A_{ij}x_j\Big\|_{\psi_2}^2 \lesssim \sum_{j=1}^n (A_{ij}x_{ij})^2 \le \|A\|_{2\to\infty}^2.$$

Thus ξ_i^2 is a sub-exponential random variable (see [Ver12, Lemma 5.9]) and

$$\|\xi_i^2\|_{\psi_1} \lesssim \|\xi_i\|_{\psi_2}^2 \lesssim \|A\|_{2\to\infty}^2. \tag{4.28}$$

Applying Bernstein's concentration inequality [Ver12, Corollary 5.17] together with

(4.27) and (4.28), we obtain

$$\mathbb{P}\left\{\sum_{i=1}^{n}\xi_i^2 \geq n\|A\|_{2\to\infty}^2 + tn\|A\|_{2\to\infty}^2\right\} \leq \exp(-ctn)$$

for all $t \geq 1$. Thus we obtained a bound on $\|\tilde{A}x\|_2^2 = \sum_{i=1}^{n}\xi_i^2$. It remains to recall (4.26) and take a union bound over $x \in \{-1,1\}^n$. It follows that the inequality

$$\|\tilde{A}\|_{\infty\to 2}^2 \leq (1+t)n\|A\|_{2\to\infty}^2$$

holds with probability at least

$$1 - 2^n\exp(-ctn) \geq 1 - \exp\left[(1-ct)n\right],$$

where $t \geq 1$ is arbitrary. Integration of these tails implies (4.25). $\square$

We will need a minor variation of Lemma 4.8 that can be applied even when some of the columns of A are removed.

Lemma 4.9 (From $2 \to \infty$ to $\infty \to 2$ for symmetric distributions)**.** *Let A be an $n \times n$ random matrix whose entries are independent, symmetric random variables. Let $J \subset [n]$ be a random subset, which is independent of the signs of the entries of A. Then*

$$\|A_J\|_{\infty\to 2} \leq C\sqrt{n}\|A_J\|_{2\to\infty}$$

with probability at least $1 - e^{-n}$.

Proof. It is quite straightforward to check this result by modifying the proof of Lemma 4.8. By the symmetry assumption, the matrix $\tilde{A} := (\varepsilon_{ij}A_{ij})$ has the same distribution as A. Conditioning on A and J leaves all randomness with the signs (ε_{ij}), as before. Then we repeat the rest of the proof of Lemma 4.8 for the submatrix A_J. In the end, we choose t to be a large absolute constant to complete the proof. $\square$

So, the only part of Lemma 4.8 that does not work for a matrix with removed columns is the symmetrization part. In the following two sections we will develop the tools to overcome the extra symmetry assumption we have to add in Lemma 4.9.

We just showed how to convert an $\infty \to 2$ bound to a $2 \to \infty$ bound for random matrices by using random signs. Alternatively, one can use random permutations for the same purpose, and obtain the following bound.

Lemma 4.10 (From $2 \to \infty$ to $\infty \to 2$). *Let A be an $n \times n$ random matrix with i.i.d. entries. Then*

$$\mathbb{E}\|A\|_{\infty\to 2} \le C\sqrt{n} \cdot \mathbb{E}\|A\|_{2\to\infty} + C\mathbb{E}\|A\boldsymbol{1}\|_2,$$

where $\boldsymbol{1} = (1, 1, \dots, 1)$ denotes the vector whose all coordinates equal 1.

Before we turn to the proof, note that the only difference between Lemmas 4.8 and 4.10 is the term $\mathbb{E}\|A\mathbf{1}\|_2$. It makes its appearance since there is no mean zero assumption on the entries. This term is usually quite innocent. Note also that (4.3) trivially implies that

$$\mathbb{E}\|A\|_{\infty\to 2} \ge \mathbb{E}\|A\mathbf{1}\|_2,$$

so we have to control this term anyway.

Proof. Let us apply a random independent permutation π_i to the elements of each row of A. The resulting matrix $\tilde{A}$ has the same distribution of A due to the i.i.d. assumption. Condition on A; the randomness now rests in the random permutations π_i only. It suffices to show that the conditional expectation satisfies

$$\mathbb{E}\|\tilde{A}\|_{\infty\to 2} \le C\sqrt{n} \cdot \|A\|_{2\to\infty} + C\|A\mathbf{1}\|_2, \tag{4.29}$$

Similarly to the proof of Lemma 4.8, we express $\|\tilde{A}x\|^2$ as a sum of independent random variables

$$\|\tilde{A}x\|_2^2 = \sum_{i=1}^n \xi_i^2 \quad \text{where} \quad \xi_i := \langle \tilde{A}_i, x\rangle = \sum_{j=1}^n A_{i,\pi_i(j)} x_j. \tag{4.30}$$

The concentration inequality for random permutations (Lemma 2.11) states that each ξ_i is a subgaussian random variable, and we have

$$\|\xi_i - \mathbb{E}\xi_i\|_{\psi_2} \lesssim \|\tilde{A}_i\|_2 \le \|A\|_{2\to\infty}.$$

Just like in the proof of Lemma 4.8, this implies that

$$\|(\xi_i - \mathbb{E}\xi_i)^2\|_{\psi_1} \lesssim \|A\|_{2\to\infty}^2.$$

Since the expectation is bounded by the ψ_1 norm (see e.g. [Ver12, Definition 5.13]), we conclude that

$$\mathbb{E}(\xi_i - \mathbb{E}\xi_i)^2 \lesssim \|(\xi_i - \mathbb{E}\xi_i)^2\|_{\psi_1} \lesssim \|A\|_{2\to\infty}^2$$

and thus

$$\mathbb{E}\sum_{i=1}^{n}(\xi_i - \mathbb{E}\xi_i)^2 \lesssim n\|A\|_{2\to\infty}^2.$$

Applying Bernstein's inequality like in Lemma 4.8, we find that

$$\mathbb{P}\left\{\sum_{i=1}^{n}(\xi_i - \mathbb{E}\xi_i)^2 \geq n\|A\|_{2\to\infty}^2 + tn\|A\|_{2\to\infty}^2\right\} \leq \exp\left[-c\min(t^2, t)n\right]$$

for all $t \geq 0$. Thus, for any $t \geq 1$ we have with probability at least $1 - \exp(-tn)$ that

$$\sum_{i=1}^{n}(\xi_i - \mathbb{E}\xi_i)^2 \leq (1+t)n\|A\|_{2\to\infty}^2. \tag{4.31}$$

From (4.30) we see that we are almost done; we just need to remove $\mathbb{E}\xi_i$ from our bound. To this end, note that

$$\|\tilde{A}x\|_2^2 = \sum_{i=1}^{n}\xi_i^2 \leq 2\sum_{i=1}^{n}(\xi_i - \mathbb{E}\xi_i)^2 + 2\sum_{i=1}^{n}(\mathbb{E}\xi_i)^2. \tag{4.32}$$

We have already bounded the first sum. As for the second one, the definition of ξ in (4.30) yields

$$\mathbb{E}\xi_i = \frac{2m-n}{n}\sum_{j=1}^{n}A_{ij} = \frac{2m-n}{n}\langle A_i, \mathbf{1}\rangle$$

where m denotes the number of ones in x_j and A_i^{T} is the i-th row of A. Thus

$$\sum_{i=1}^{n}(\mathbb{E}\xi_i)^2 = \Big(\frac{2m-n}{n}\Big)^2\sum_{i=1}^{n}\langle A_i, \mathbf{1}\rangle^2 \leq \|A\mathbf{1}\|_2^2.$$

We substitute this and (4.31) into (4.32) and obtain that for any $t \geq 1$,

$$\|\tilde{A}x\|_2^2 \leq 2(1+t)n\|A\|_{2\to\infty}^2 + 2\|A\mathbf{1}\|_2^2$$

with probability at least $1 - \exp(-tn)$.

It remains to recall (4.26) and take a union bound over $x \in \{-1, 1\}^n$. It follows that the inequality

$$\|\tilde{A}\|_{\infty\to 2}^2 \leq 2(1+t)n\|A\|_{2\to\infty}^2 + 2\|A\mathbf{1}\|_2^2 \tag{4.33}$$

holds with probability at least

$$1 - 2^n \exp(-ctn) \geq 1 - \exp\left[(1-ct)n\right]$$

where $t \geq 1$ is arbitrary. Integration of these tails implies (4.29). $\square$

It is worthwhile to mention a high-probability version of Lemma 4.10.

Lemma 4.11 (From $2 \to \infty$ to $\infty \to 2$ with high probability)**.** *Let A be an $n \times n$ random matrix with i.i.d. entries. Then with probability at least $1 - e^{-n}$ we have*

$$\|A\|_{\infty\to 2} \leq C\sqrt{n} \cdot \mathbb{E}\|A\|_{2\to\infty} + C\mathbb{E}\|A\boldsymbol{1}\|_2,$$

where $\boldsymbol{1} = (1, 1, \ldots, 1)$ denotes the vector whose all coordinates equal 1.

Proof. At the end of the proof of Lemma 4.10, we obtained inequality (4.33) which states (for large constant t) that

$$\|\tilde{A}\|_{\infty\to 2} \leq C\sqrt{n} \cdot \|A\|_{2\to\infty} + C\|A\mathbf{1}\|_2$$

with probability at least $1 - e^{-n}$. Note that

$$\|A\|_{2\to\infty} = \|\tilde{A}\|_{2\to\infty} \quad \text{and} \quad \|A\mathbf{1}\|_2 = \|\tilde{A}\mathbf{1}\|_2$$

deterministically. Indeed, it is easy to check that permutations of the elements of the rows of A do not affect these two quantities. It follows that

$$\|\tilde{A}\|_{\infty\to 2} \leq C\sqrt{n} \cdot \|\tilde{A}\|_{2\to\infty} + C\|\tilde{A}\mathbf{1}\|_2$$

with probability at least $1 - e^{-n}$. It remains to note that $\tilde{A}$ has the same distribution as A. $\square$

Recall from Section 4.1.1 that ideally, we would want

$$\|A\|_{2\to\infty} \lesssim \sqrt{n} \quad \text{and} \quad \|A\|_{\infty\to 2} \lesssim n$$

with high probability. But this is too good to be true in our situation, where we assume only two moments for the entries of A. Nevertheless, we will now show that these bounds still hold, albeit with exponentially small probability.

Lemma 4.12 ($2 \to \infty$ and $\infty \to 2$ norms with tiny probability)**.** *Let A be an $n \times n$ random matrix whose entries are i.i.d. random variables with mean zero and at most unit variance. Let $\delta \in (0, 1/2)$. Then*

$$\|A\|_{2\to\infty} \le 2\delta^{-1}\sqrt{n} \quad \textit{and} \quad \|A\|_{\infty\to 2} \le C\delta^{-1}n \tag{4.34}$$

with probability at least $\frac{1}{2}\exp(-\delta^2 n)$.

Proof. We will first bound below the probability of the event

$$\mathcal{E} := \Big\{\|A\|_{2\to\infty} \le 2\delta^{-1}\sqrt{n} \text{ and } \|\tilde{A}\mathbf{1}\|_2 \le 2\delta^{-1}n\Big\}$$

and then use Lemma 4.11 to control $\|A\|_{\infty\to 2}$.

Recall from (4.2) that

$$\|A\|_{2\to\infty} = \max_{i\in[n]} \|A_i\|_2 \quad \text{and} \quad \|\tilde{A}\mathbf{1}\|_2^2 = \sum_{i=1}^{n} \langle A_i, \mathbf{1}\rangle^2$$

where A_i^T denote the rows of A. Thus $\mathcal{E} \subset \bigcap_{i=1}^n \mathcal{E}_i$ where

$$\mathcal{E}_i := \big\{\|A_i\|_2 \le 2\delta^{-1}\sqrt{n} \text{ and } |\langle A_i, \mathbf{1}\rangle| \le 2\delta^{-1}\sqrt{n}\big\}$$

are independent events. This reduces the problem to bounding the probability of each event $\mathcal{E}_i$ below.

The assumptions on the entries of A imply that

$$\mathbb{E}\|A_i\|_2^2 \le n \quad \text{and} \quad \mathbb{E}\langle A_i, \mathbf{1}\rangle^2 \le n.$$

Using Chebyshev's inequality, we see that

$$\mathbb{P}\left\{\|A_i\|_2 > 2\delta^{-1}\sqrt{n}\right\} \le \frac{\delta^2}{4} \quad \text{and} \quad \mathbb{P}\left\{|\langle A_i, \mathbf{1}\rangle| > 2\delta^{-1}\sqrt{n}\right\} \le \frac{\delta^2}{4}.$$

Then a union bound yields

$$\mathbb{P}(\mathcal{E}_i) \ge 1 - \frac{\delta^2}{2}.$$

By independence of the events $\mathcal{E}_i$, this implies

$$\mathbb{P}(\mathcal{E}) \ge \left(1 - \frac{\delta^2}{2}\right)^n \ge \exp(-\delta^2 n).$$

Next we apply Lemma 4.11, which states that the event

$$\mathcal{F} := \left\{ \|A\|_{\infty\to 2} \le C\sqrt{n} \cdot \mathbb{E}\|A\|_{2\to\infty} + C\mathbb{E}\|A\mathbf{1}\|_2 \right\}$$

is likely:

$$\mathbb{P}(\mathcal{F}) \ge 1 - \exp(-n).$$

It follows that

$$\mathbb{P}(\mathcal{E} \cap \mathcal{F}) \ge \exp(-\delta^2 n) - \exp(-n) \ge \frac{1}{2}\exp(-\delta^2 n).$$

It remains to note that by definition of $\mathcal{E}$ and $\mathcal{F}$, the event $\mathcal{E} \cap \mathcal{F}$ implies the inequalities in (4.34). □

In the previous section, we were able to prove the optimal bounds

$$\|A\|_{2\to\infty} \lesssim \sqrt{n} \quad \text{and} \quad \|A\|_{\infty\to 2} \lesssim n$$

for a random matrix A, but they only hold with exponentially small probability. We claim that the probability of success can be increased to almost 1 if we are allowed to remove a few columns of A. We already proved this fact for the $2 \to \infty$ norm in Lemma 4.7. It is time to handle the $\infty \to 2$ norm.

Theorem 4.13 (Bounding $\infty \to 2$ norm by removing a few columns)**.** *Consider an $n \times n$ random matrix A with i.i.d. entries A_{ij} which have mean zero and at most unit variance and satisfy* (4.22)*. Let $\varepsilon \in (0, 1/2]$. Then with probability at least $1 - 2\exp(-\varepsilon n)$, there exists a subset $J \in [n]$ with cardinality $|J| \le \varepsilon n$ such that*

$$\|A_{J^c}\|_{\infty\to 2} \le C\sqrt{\ln \varepsilon^{-1}} \cdot n.$$

Proof. **Step 1: Defining the two key events.** We will be interested in the two key events that suitably control the $2 \to \infty$ and $\infty \to 2$ norms of a random matrix. Thus, for a random matrix B and numbers $r, K \ge 0$, we define

$$\mathcal{E}_{2\to\infty}(B, r, K) := \left\{ \exists J,\ |J| \le r\varepsilon n :\ \|B_{J^c}\|_{2\to\infty} \le K\sqrt{\ln \varepsilon^{-1}} \cdot \sqrt{n} \right\},$$
$$\mathcal{E}_{\infty\to 2}(B, r, K) := \left\{ \exists J,\ |J| \le r\varepsilon n :\ \|B_{J^c}\|_{\infty\to 2} \le K\sqrt{\ln \varepsilon^{-1}} \cdot n \right\}.$$

In terms of these events, we want to show that

$$\mathbb{P}\left(\mathcal{E}_{\infty\to 2}(A,1,C)^c\right) \le 2\exp(-\varepsilon n),$$

while Lemma 4.7 can be stated as

$$\mathbb{P}\left(\mathcal{E}_{2\to\infty}(A,1,C')\right) \ge 1-\exp(-\varepsilon n).$$

for some absolute constant C'. Since the latter event is so likely, intersecting with it would not cause much harm. Indeed, we will show that the bad event

$$\mathcal{B} := \mathcal{E}_{2\to\infty}(A,1,C') \cap \mathcal{E}_{\infty\to 2}(A,1,C)^c$$

satisfies

$$\mathbb{P}(\mathcal{B}) \le \exp(-n/2). \tag{4.35}$$

This would finish the proof, since we would then have

$$\mathbb{P}\left(\mathcal{E}_{\infty\to 2}(A,1,C)^c\right) \le \exp(-n/2)+\exp(-\varepsilon n) \le 2\exp(-\varepsilon n)$$

as required.

Step 2: Symmetrization. As an intermediate step, let us bound the probability of a symmetrized version of $\mathcal{B}$, namely the event

$$\tilde{\mathcal{B}} := \mathcal{E}_{2\to\infty}(\tilde{A},1,2C') \cap \mathcal{E}_{\infty\to 2}(\tilde{A},1,C/2)^c$$

where

$$\tilde{A} := A - A'$$

and A' is an independent copy of the random matrix A. We claim that

$$\mathbb{P}(\tilde{\mathcal{B}}) \le \exp(-n). \tag{4.36}$$

To prove this claim, choose a subset J, $|J| \le \varepsilon n$, that minimizes $\|\tilde{A}_{J^c}\|_{2\to\infty}$. Recall from (4.2) that the $2 \to \infty$ norm of a matrix is determined by the Euclidean norms of the columns and thus does not depend on the signs of the matrix elements. Thus J is independent of the signs of the elements of $\tilde{A}$. This makes it possible to use Lemma 4.9 for the matrix $\tilde{A}$ and the random set J^c. It gives

$$\|\tilde{A}_{J^c}\|_{\infty\to 2} \lesssim \sqrt{n}\|\tilde{A}_{J^c}\|_{2\to\infty} \tag{4.37}$$

with probability at least $1-\exp(-n)$.

Then, turning to $\tilde{\mathcal{B}}$, we can bound its probability as follows:

$$\mathbb{P}(\tilde{\mathcal{B}}) \le \mathbb{P}(\tilde{\mathcal{B}} \text{ and } (4.37)) + \exp(-n).$$

To prove the claim, it remains to check that $\tilde{\mathcal{B}}$ and (4.37) can not hold together. Assume they do; then

$$\|\tilde{A}_{J^c}\|_{\infty\to 2} \lesssim \sqrt{n}\cdot 2C'\sqrt{\ln\varepsilon^{-1}}\sqrt{n} \lesssim \sqrt{\ln\varepsilon^{-1}}\cdot n,$$

which contradicts the event $\mathcal{E}_{\infty\to 2}(\tilde{A},1,C/2)^c$ in the definition of $\tilde{\mathcal{B}}$ for a suitably chosen constant C. This completes the proof of the claim (4.36).

Step 3. Using the small-probability bounds. The last piece of information we will use is the conclusion of Lemma 4.12 for $\delta := 1/(2\ln\varepsilon^{-1})$. It states that the good event

$$\mathcal{G} := \mathcal{E}_{2\to\infty}(A',0,C') \cap \mathcal{E}_{\infty\to 2}(A',0,C/2)$$

is likely to happen:

$$\mathbb{P}(\mathcal{G}) \ge \frac{1}{2}\exp\Big(-\frac{n}{4\ln\varepsilon^{-1}}\Big). \tag{4.38}$$

Note in passing that there is no guarantee that this statement would hold for the same constants C and C' as we chose in the definition of $\mathcal{B}$ above. However, we can make this happen by adjusting these constants upwards as necessary. The reader can easily check both (4.36) and (4.38) would still hold after such an adjustment.

We claim that

$$\mathcal{B}\cap\mathcal{G} \subset \tilde{\mathcal{B}}. \tag{4.39}$$

To see this, recall that each of $\mathcal{B}$, $\mathcal{G}$ and $\tilde{\mathcal{B}}$ is defined an an intersection of two events, one controlling $2\to\infty$ norm and the other, $\infty\to 2$ norm. Thus it suffices to check the inclusion for each of these two parts separately. Namely, the claim (4.39) would follow at once if we show that

$$\mathcal{E}_{2\to\infty}(A,1,C') \cap \mathcal{E}_{2\to\infty}(A',0,C') \subset \mathcal{E}_{2\to\infty}(\tilde{A},1,2C') \quad \text{and}$$
$$\mathcal{E}_{\infty\to 2}(A,1,C)^c \cap \mathcal{E}_{\infty\to 2}(A',0,C/2) \subset \mathcal{E}_{\infty\to 2}(\tilde{A},1,C/2)^c.$$

Both these inclusions are straightforward to check from the definitions of the events $\mathcal{E}_{2\to\infty}$

and $\mathcal{E}_{\infty\to 2}$, remembering that $\tilde{A} = A - A'$ and using triangle triangle inequality. This verifies the claim (4.39).

The event $\mathcal{B}$ is determined by A, and $\mathcal{G}$ is determined by A' only. Thus $\mathcal{B}$ and $\mathcal{G}$ are independent, and (4.39) gives

$$\mathbb{P}(\mathcal{B})\,\mathbb{P}(\mathcal{G}) = \mathbb{P}(\mathcal{B} \cap \mathcal{G}) \le \mathcal{P}(\tilde{\mathcal{B}}).$$

Thus, using (4.36) and (4.38), we conclude that

$$\mathbb{P}(\mathcal{B}) \le \mathbb{P}(\tilde{\mathcal{B}})/\mathbb{P}(\mathcal{G}) \le 2\exp\Big(-n + \frac{n}{4\ln\varepsilon^{-1}}\Big) \le \exp(-n/2).$$

We have shown (4.35) and thus have completed the proof of the theorem. □

From $\infty \to 2$ norm to the operator norm

In Theorem 4.13, we gave an optimal $O(n)$ bound for the $\infty \to 2$ norm of a random matrix with few removed columns. We will now convert this into an optimal $O(\sqrt{n})$ bound for the operator norm. This can be done by applying a form of Grothendieck-Pietsch theorem (see [LT13, Proposition 15.11]), which has been used recently in [LLV17, section 3.2] in a similar context.

Theorem 4.14 (Grothendieck-Pietsch)**.** *Let B be a $k \times m$ real matrix and $\delta > 0$. Then there exists $J \subset [m]$ with $|J| \le \delta m$ such that*

$$\|B_{J^c}\| \le \frac{2\|B\|_{\infty\to 2}}{\sqrt{\delta m}}.$$

Applying Theorem 4.13 followed by Grothendieck-Pietsch theorem, we obtain the following result.

Lemma 4.15 (Bounding the operator norm by removing a few columns)**.** *Consider an $n\times n$ random matrix A with i.i.d. entries A_{ij} which have mean zero and at most unit variance and satisfy* (4.22)*. Let $\varepsilon \in (0,1]$. Then with probability at least $1 - 2\exp(-\varepsilon n/2)$, there exists a subset $J \in [n]$ with cardinality $|J| \le \varepsilon n$ such that*

$$\|A_{J^c}\| \le C\sqrt{\frac{\ln\varepsilon^{-1}}{\varepsilon}} \cdot \sqrt{n}.$$

Proof. Apply Theorem 4.13 for $\varepsilon/2$ instead of ε. We obtain a subset of columns $J_1 \subset [n]$,

$|J_1| \le \varepsilon n/2$, which satistfies

$$\|A_{J_1^c}\|_{\infty\to 2} \le C\sqrt{\ln \varepsilon^{-1}} \cdot n \tag{4.40}$$

with probability at least $1 - 2\exp(-\varepsilon n/2)$.

Next apply Grothendieck-Pietsch Theorem 4.14 for the matrix $A_{J_1^c}$ and for $\delta = \varepsilon/2$. We obtain a further subset $J_2 \subset J_1^c$, $|J_2| \le \delta |J_1^c| \le \varepsilon n/2$, such that the removal of columns in both $J := J_1 \cup J_2$ leads to

$$\|A_{J^c}\| \le \frac{2\|A_{J_1^c}\|_{\infty\to 2}}{\sqrt{\delta |J_1^c|}} \lesssim C\sqrt{\frac{\ln \varepsilon^{-1}}{\varepsilon}} \cdot \sqrt{n}.$$

In the last inequality, we used the bound (4.40) and that $\delta = \varepsilon/2$ and $|J_1^c| \ge n - \varepsilon n/2 \ge n/2$. The proof is complete. □

We are ready to prove a partial case of Theorem 4.1, for the matrices whose entries are $O(\sqrt{n})$. It follows by applying Lemma 4.15 for A and A^T separately, and then superposing the results.

Proposition 4.16 (Bounded entries)**.** *Consider an $n \times n$ random matrix A with i.i.d. entries A_{ij} which have mean zero and at most unit variance and satisfy* (4.22)*. Let $\varepsilon \in (0, 1]$. Then with probability at least $1 - 4\exp(-\varepsilon n/2)$, there exists an $\varepsilon n \times \varepsilon n$ submatrix of A such that replacing all of its entries with zero leads to a well-bounded matrix $\tilde{A}$:*

$$\|\tilde{A}\| \le C\sqrt{\frac{\ln \varepsilon^{-1}}{\varepsilon}} \cdot \sqrt{n}.$$

Proof. Apply Lemma 4.15 for A and A^T. We obtain that with probability at least $1 - 4\exp(-\varepsilon n/2)$, there exists sets I and J with at most εn indices in each, and such that

$$\|A_{[n]\times J^c}\| \lesssim \sqrt{\frac{\ln \varepsilon^{-1}}{\varepsilon}} \cdot \sqrt{n} \quad \text{and} \quad \|A_{I^c\times [n]}\| \lesssim \sqrt{\frac{\ln \varepsilon^{-1}}{\varepsilon}} \cdot \sqrt{n}. \tag{4.41}$$

We claim that $\tilde{A} := A_{(I\times J)^c}$ satisfies the conclusion of the proposition. The support of this matrix, $(I \times J)^c$, is a disjoint union of two sets, $[n] \times J^c$ and $I^c \times J$. Then, using the triangle inequality, we have

$$\|A_{(I\times J)^c}\| \le \|A_{[n]\times J^c}\| + \|A_{I^c\times J}\|.$$

We already controlled the first term in (4.41). As for the second term, since adding columns can only increase the operator norm, we have $\|A_{I^c\times J}\| \le \|A_{I^c\times [n]}\|$, which we also bounded in (4.41). The proof is complete. □

4.2.3 Controlling the moderately large entries: Bernoulli matrices

In the previous section, we proved a partial case of Theorem 4.1 that controls relatively small entries of A, those of the order $O(\sqrt{n})$. Larger entries will be controlled in this section.

The following general lemma will help us analyze the patterns such large entries can form.

Lemma 4.17 (Bernoulli random matrix)**.** *Let B be an $n \times n$ random matrix whose entries are independent Bernoulli random variables with mean p. Let $\varepsilon \in (0, 1/2]$. Consider the rows of B with more than $21pn + 2\ln\varepsilon^{-1}$ ones. Then with probability $1 - \exp(-\varepsilon n/2)$, these rows have at most εn ones altogether.*

To see the connection to our original problem, we will later choose the entries of B to be the indicators of the large entries of A.

Proof. Let S_i denote the number of ones in the i-th row of B. Then $\mathbb{E}S_i = pn$. A standard application of Chernoff's inequality shows that

$$\mathbb{P}\left\{S_i > t\right\} \le e^{-2t} \quad \text{for } t \ge 21pn. \tag{4.42}$$

Let $K \ge 21pn$ be a number to be chosen later. (We will eventually choose K as $21pn + 2\ln\varepsilon^{-1}$ as in the statement of the lemma.) Define the random variables

$$X_i := S_i \mathbb{1}_{\{S_i > K\}}.$$

The quantity of interest is the total number of ones in the heavy rows, and it equals $\sum_{i=1}^n X_i$. To control this sum of independent random variables, we can use the standard Bernstein's trick (commonly called Chernoff's bound), where we use Markov's inequality after exponentiation. We obtain

$$\mathbb{P}\left\{\sum_{i=1}^n X_i > \varepsilon n\right\} \le e^{-\varepsilon n}\mathbb{E}\exp\Big(\sum_{i=1}^n X_i\Big) = \Big[e^{-\varepsilon}\mathbb{E}e^{X_1}\Big]^n, \tag{4.43}$$

where the last equality follows by independence and identical distribution. Now, by definition of X_1 we have

$$
\begin{aligned}
\mathbb{E}e^{X_1} &= \mathbb{E}e^{X_1}\mathbb{1}_{\{X_1=0\}} + \mathbb{E}e^{X_1}\mathbf{1}_{\{X_1\neq 0\}} \le 1 + \mathbb{E}e^{S_1}\mathbb{1}_{\{S_1>K\}} \\
&= 1 + \int_{e^K}^{\infty} \mathbb{P}\left\{e^{S_1} > u\right\} du \\
&= 1 + \int_{K}^{\infty} \mathbb{P}\left\{S_1 > t\right\} e^t \, dt \quad \text{(by a change of variables)} \\
&\le 1 + \int_{K}^{\infty} e^{-2t} e^t \, dt \quad \text{(using (4.42) for } t \ge K \ge 21pn\text{)} \\
&= 1 + e^{-K} \le \exp(e^{-K}).
\end{aligned}
$$

Substituting this bound into (4.43), we conclude that

$$
\mathbb{P}\left\{\sum_{i=1}^{n} X_i > \varepsilon n\right\} \le \exp\left[(-\varepsilon + e^{-K})n\right] \le \exp(-\varepsilon n/2),
$$

if we choose K so that $e^{-K} \le \varepsilon/2$. To finish the proof, recall that our argument works if K satisfies the two conditions: $K \ge 21pn$ and $e^{-K} \le \varepsilon/2$. We thus choose $K := 21pn + 2\ln\varepsilon^{-1}$ and complete the proof. □

Corollary 4.18 (Bernoulli random matrix)**.** *Let B be an $n \times n$ random matrix whose entries are independent Bernoulli random variables with mean p. Let $\varepsilon \in (0, 1]$. Then with probability at least $1 - 2\exp(-\varepsilon n/4)$, there exists an $\varepsilon n \times \varepsilon n$ submatrix of B such that replacing all of its entries with zero leads to a matrix $\tilde{B}$ whose rows and columns have at most $21pn + 4\ln\varepsilon^{-1}$ ones each.*

Proof. Apply Lemma 4.17 for B and B^{T} with $\varepsilon/2$ instead of ε, and take the intersection of the two good events. With the required probability, we obtain a set of εn bad entries of B whose removal makes all rows and columns of B contain at most $21pn + 2\ln\varepsilon^{-1}$ ones. It remains to note that these εn entries can be trivially placed in some $\varepsilon n \times \varepsilon n$ submatrix of B, and deletion of the whole $\varepsilon n \times \varepsilon n$ submatrix can only decrease the number of non-zero elements in the rows and columns of the residual part. □

Remark 4.19 (Random graphs). It is not difficult to obtain a version of Corollary 4.18 for symmetric random matrices. This version can be interpreted as a statement about Erdős-Rényi random graphs $G(n, p)$, with B playing the role of the adjacency matrix. It states that with high probability, one can make all degrees of a $G(n, p)$ random graph bounded by $O(pn + \ln\varepsilon^{-1})$ after removing the internal edges from a sub-graph with εn vertices.

We will use Corollary 4.18 to deduce Theorem 4.1 for matrices with moderately large entries. Namely, we assume here that all entries of A satisfy

$$A_{ij} = 0 \quad \text{or} \quad \frac{\sqrt{n}}{2} \leq |A_{ij}| \leq \frac{5\sqrt{n}}{\sqrt{\varepsilon}}. \tag{4.44}$$

Proposition 4.20 (Moderately large entries)**.** *Consider an $n \times n$ random matrix A with i.i.d. entries which satisfy $\mathbb{E}A_{ij}^2 \leq 1$ and (4.44). Let $\varepsilon \in (0, 1/2]$. Then with probability at least $1 - 2\exp(-\varepsilon n/4)$, there exists an $\varepsilon n \times \varepsilon n$ submatrix of A such that replacing all of its entries with zero leads to a well-bounded matrix $\tilde{A}$:*

$$\|\tilde{A}\| \leq \frac{C \ln \varepsilon^{-1}}{\sqrt{\varepsilon}} \cdot \sqrt{n}. \tag{4.45}$$

Proof. Consider the matrix B whose elements are indicators of moderately large entries of A, i.e.

$$B_{ij} := \mathbb{1}_{\{A_{ij} \neq 0\}}.$$

Then B_{ij} are i.i.d. Bernoulli random variables with mean

$$p := \mathbb{E}B_{ij} = \mathbb{P}\left\{A_{ij} \neq 0\right\} \leq \mathbb{P}\left\{|A_{ij}| \geq \frac{\sqrt{n}}{2}\right\} \leq \frac{4}{n}. \tag{4.46}$$

(In the last inequality, we used Chebyshev's inequality and the assumption $\mathbb{E}A_{ij}^2 \leq 1$.) Corollary 4.18 applied to B gives us an $\varepsilon n \times \varepsilon n$ submatrix of A such that the number of non-zero elements in every row and column of $\tilde{A}$ (obtained by zeroing out the elements of A outside that submatrix) is bounded by

$$21pn + 4\ln\varepsilon^{-1} \lesssim \ln\varepsilon^{-1}, \tag{4.47}$$

where we used (4.46) in the last bound.

Moreover, assumption (4.44) shows that all entries of $\tilde{A}$ are bounded in absolute value by $5\sqrt{n}/\sqrt{\varepsilon}$. This and (4.47) imply that the ℓ_1 norm of all rows $\tilde{A}_i$ and columns $\tilde{A}^j$ can be bounded as follows:

$$\max_{i,j}\left(\|\tilde{A}_i\|_1,\, \|\tilde{A}^j\|_1\right) \lesssim \frac{\sqrt{n}}{\sqrt{\varepsilon}} \cdot \ln\varepsilon^{-1}.$$

Applying Lemma 2.19 leads to (4.45). $\square$

4.2.4 Controlling the large entries and proof of Theorem 4.1

Finally, we will need to prove Theorem 4.1 for very large entries – now we assume that all entries of A satisfy

$$A_{ij} = 0 \quad \text{or} \quad |A_{ij}| > \frac{5\sqrt{n}}{\sqrt{\varepsilon}}. \tag{4.48}$$

There are typically very few such entries, as the following simple result shows.

Lemma 4.21 (Few very large entries)**.** *Consider an $n \times n$ random matrix A with i.i.d. entries which satisfy $\mathbb{E}A_{ij}^2 \le 1$ and* (4.48)*. Let $\varepsilon \in (0, 1/2]$. Then with probability at least $1 - \exp(-\varepsilon n)$, the matrix A has at most εn non-zero entries.*

Proof. Using Chebyshev's inequality and the assumption that $\mathbb{E}A_{ij}^2 \le 1$, we see that the probability that a given entry is nonzero is

$$\mathbb{P}\left\{A_{ij} \neq 0\right\} \le \mathbb{P}\left\{|A_{ij}| > \frac{5\sqrt{n}}{\sqrt{\varepsilon}}\right\} \le \frac{\varepsilon}{25n}.$$

Thus the expected number of non-zero entries in A is at most $\varepsilon n/25$. A standard application of Chernoff's inequality gives

$$\mathbb{P}\left\{A \text{ has more than } \varepsilon n \text{ nonzero entries}\right\} \le e^{-\varepsilon n}.$$

The proof is complete. $\square$

Since a set of εn indices can be always placed in an $\varepsilon n \times \varepsilon n$ submatrix, we can state Lemma 4.21 as follows.

Corollary 4.22 (Very large entries)**.** *Consider an $n \times n$ random matrix A with i.i.d. entries which satisfy $\mathbb{E}A_{ij}^2 \le 1$ and* (4.48)*. Let $\varepsilon \in (0, 1/2]$. Then with probability at least $1 - \exp(-\varepsilon n)$, all non-zero entries of A are contained in an $\varepsilon n \times \varepsilon n$ submatrix.*

We are going to assemble Proposition 4.16 for the bounded entries of A, Proposition 4.20 for moderately large entries, and Corollary 4.22 for very large entries. The $\varepsilon n \times \varepsilon n$ sub-matrices that appear in these results are possibly different. The following simple lemma will help us to combine them into one.

Lemma 4.23. *Let B be a matrix. Zeroing out any submatrix of B cannot increase the operator norm more than twice.*

Proof. Let $\tilde{B}$ denotes the matrix obtained by zeroing out a submatrix of B spanned by the index set $I \times J$. Triangle inequality gives

$$\|\tilde{B}\| \le \|B_{I \times J^c}\| + \|B_{I^c \times [n]}\| \le \|B\| + \|B\|.$$

The last inequality follows from the fact that zeroing out any subset of rows or columns cannot increase the operator norm. □

Proof of Theorem 4.1. Decompose A into a sum of three $n \times n$ matrices with disjoint support,

$$A = B + M + L, \tag{4.49}$$

where B contains bounded entries of A – those that satisfy $|A_{ij}| \le \sqrt{n}/2$, the matrix M contains moderately large entries – those for which $\sqrt{n}/2 < |A_{ij}| \le 5\sqrt{n/\varepsilon}$, and L contains large entries – those satisfying $|A_{ij}| > 5\sqrt{n/\varepsilon}$.

To bound B, let us subtract the mean and first bound

$$G := B - \mathbb{E}B.$$

The entries of this matrix have zero mean and satisfy

$$\mathbb{E}G_{ij}^2 = \mathrm{Var}(B_{ij}) \le \mathbb{E}B_{ij}^2 \le \mathbb{E}A_{ij}^2 = 1$$

(where we used the moment assumption) and

$$\|G_{ij}\|_\infty = \|B_{ij} - \mathbb{E}B_{ij}\|_\infty \le 2\|B_{ij}\|_\infty \le \sqrt{n}.$$

Thus we can apply Proposition 4.16 for $0.5G$. It says that with probability at least $1 - 4\exp(-\varepsilon n/2)$, the removal of a certain $\varepsilon n \times \varepsilon n$ submatrix of G leads to a well-bounded matrix $\tilde{G}$, i.e.

$$\|\tilde{G}\| \lesssim \sqrt{\frac{\ln \varepsilon^{-1}}{\varepsilon}} \cdot \sqrt{n}.$$

Next, we bound $\mathbb{E}B$, a matrix whose entries are the same. Thus

$$\begin{aligned}
\|\mathbb{E}B\| &= n\big|\mathbb{E}B_{ij}\big| = n\big|\mathbb{E}(A_{ij} - B_{ij})\big| \quad \text{(since } \mathbb{E}A_{ij} = 0\text{)} \\
&= n\big|\mathbb{E}A_{ij}\mathbb{1}_{|A_{ij}|>\sqrt{n}/2}\big| \quad \text{(by definition of } A_1\text{)} \\
&\le n\big(\mathbb{E}A_{ij}^2\big)^{1/2}\big(\mathbb{P}\left\{|A_{ij}| > \sqrt{n}/2\right\}\big)^{1/2} \quad \text{(by Cauchy-Schwarz inequality)} \\
&\le n \cdot 1 \cdot 2/\sqrt{n} \le 2\sqrt{n} \quad \text{(by Chebyshev's inequality).}
\end{aligned} \tag{4.50}$$

To bound M and L, note that

$$\mathbb{E}M_{ij}^2 \le \mathbb{E}A_{ij}^2 = 1 \text{ and } \mathbb{E}L_{ij}^2 \le \mathbb{E}A_{ij}^2 = 1.$$

Thus, Proposition 4.20 can be applied to M: with probability $1 - \exp(-\varepsilon n/4)$ there exist a $\varepsilon n \times \varepsilon n$ submatrix of M, the removal of which leads to a matrix $\tilde{M}$, such that

$$\|\tilde{M}\| \leq \frac{C \ln \varepsilon^{-1}}{\sqrt{\varepsilon}} \cdot \sqrt{n}.$$

Proposition 4.22 can be applied to L. We can interpret its conclusion as follows: with probability $1 - 2\exp(-\varepsilon n)$ there exist a $\varepsilon n \times \varepsilon n$ submatrix of L, the removal of which leads to an identically zero matrix $\tilde{L}$, i.e.

$$\|\tilde{L}\| = 0.$$

Now, let us embed the three $\varepsilon n \times \varepsilon n$ submatrices of A we just constructed into one $3\varepsilon n \times 3\varepsilon n$ submatrix and zero out this whole submatrix. By Lemma 4.23, the norms of $\tilde{G}$, $\mathbb{E}B$, $\tilde{M}$ and $\tilde{L}$ will not increase more than twice as a result of this operation. Taking intersection of the three good events, we conclude that with the probability at least

$$1 - 4\exp(-\varepsilon n/2) - 2\exp(-\varepsilon n/4) - \exp(-\varepsilon n) \geq 1 - 7\exp(-\varepsilon n/4)$$

there exists an $3\varepsilon n \times 3\varepsilon n$ submatrix of A such that replacing all of its entries by zero leads to a well-bounded matrix $\tilde{A}$:

$$\|\tilde{A}\| \lesssim \sqrt{\frac{\ln \varepsilon^{-1}}{\varepsilon}} \cdot \sqrt{n} + 2\sqrt{n} + \frac{\ln \varepsilon^{-1}}{\sqrt{\varepsilon}} \cdot \sqrt{n} + 0 \lesssim \frac{\ln \varepsilon^{-1}}{\sqrt{\varepsilon}} \cdot \sqrt{n}.$$

This proves the conclusion of Theorem 4.1 with 3ε instead of ε, where $\varepsilon \in (0, 1/2]$ is arbitrary. By rescaling, Theorem 4.1 holds also as originally stated. This concludes the proof of Theorem 4.1. □

4.3 Global problem

In this section we prove Theorem 4.3, which states that either nonzero mean or infinite second moment make it impossible to repair the matrix norm by removing a small submatrix. We will first prove a non-asymptotic version version of this result. Once this is done, an application of Borel-Cantelli Lemma will quickly yield Theorem 4.3.

Proposition 4.24 (Global problem: non-asymptotic regime)**.** *Consider an $n \times n$ random matrix A whose entries are i.i.d. random variables that have either nonzero mean or infinite second moment, and let $\varepsilon \in (0, 1)$. Then, for any $M > 0$ there exists n_0 that may depend only on ε, M and the distribution of the entries, and such that for any $n > n_0$ the following* *event holds with probability at least* $1-e^{-n}$*: every* $(1-\varepsilon)n \times (1-\varepsilon)n$ *submatrix* A' *of* A *satisfies*

$$\|A'\| \geq M\sqrt{n}.$$

Before we prove this proposition, let us pause to see its connection to the matrix $\tilde{A}_n$ of Theorem 4.3. Proposition 4.24 yields that this matrix satisfies

$$\|\tilde{A}_n\| \geq M\sqrt{n}.$$

Indeed, modifying an $\varepsilon n \times \varepsilon n$ submatrix always leaves some $(1-\varepsilon)n \times (1-\varepsilon)n$ submatrix A' intact, so we can apply Proposition 4.24 for that submatrix.

4.3.1 Infinite second moment

Here we will prove the part of Proposition 4.24 about infinite second moment; the case of nonzero mean will be treated in Section 4.3.2. Let us start with the following lemma which will help us treat a fixed submatrix.

Lemma 4.25. *Consider an* $m \times m$ *random matrix* B *whose entries are i.i.d. random variables with infinite second moment. Then, for any* $M > 0$ *there exists* m_0 *that may depend only on* M *and the distribution of the entries, and such that for any* $m > m_0$ *we have*

$$\|B\| \geq M\sqrt{m}$$

with probability at least $1-\exp(-M^2 m)$.

Proof. By assumption, we have $\mathbb{E}B_{ij}^2 = \infty$. Therefore, for any $M > 0$ one can find a truncation level K that depends only on M and the distribution, and such that the truncated random variables

$$\bar{B}_{ij} := B_{ij}\mathbb{1}_{|B_{ij}|\leq K} \quad \text{satisfy} \quad \mathbb{E}\bar{B}_{ij}^2 \geq 2M^2. \tag{4.51}$$

(This follows easily from Lebesgue's monotone convergence theorem.)

Consider the matrix $\bar{B}$ with entries $\bar{B}_{ij}$. We have

$$\|B\| \geq \frac{1}{\sqrt{m}}\|B\|_F \geq \frac{1}{\sqrt{m}}\|\bar{B}\|_F.$$

Then we bound the failure probability as follows:

$$\begin{aligned}\mathbb{P}\left\{\|B\| < M\sqrt{m}\right\} \le \mathbb{P}\left\{\|\bar{B}\|_F < Mm\right\} = \mathbb{P}\left\{\sum_{i,j=1}^{m} \bar{B}_{ij}^2 < M^2m^2\right\} \\ \le \mathbb{P}\left\{\sum_{i,j=1}^{m} (\bar{B}_{ij}^2 - \mathbb{E}\bar{B}_{ij}^2) < -M^2m^2\right\}\end{aligned}$$

where we used (4.51) in the last step.

Apply Hoeffding's inequality for the random variables $\bar{B}_{ij}^2$ and use that they are bounded by K^2 by construction. The probability above gets bounded by

$$\exp\Big(-\frac{M^4m^2}{2K^2}\Big).$$

If $m > 2K^2/M^2 = m_0$, this probability can be further bounded by $\exp(-M^2m)$, as claimed. □

Proof of Proposition 4.24 for infinite second moment. We can assume without loss of generality that M is large enough depending on ε. (Indeed, once the conclusion of the proposition holds for one value of M it automatically holds for all smaller values.)

Apply Lemma 4.25 for an $m \times m$ matrix A_n' with $m = (1-\varepsilon)n$, and then take a union bound over all $\binom{n}{m}^2$ possible choices of such submatrices. It follows that the conclusion of Proposition 4.24 holds with probability at least

$$1 - \binom{n}{m}^2 \exp(-M^2m).$$

By Stirling's approximation, we have $\binom{n}{m} \le (en/m)^m$. Using this and substituting $m = (1-\varepsilon)n$, we bound the probability below by

$$1 - \exp\Big[\Big(2\log\frac{e}{1-\varepsilon} - M^2\Big)(1-\varepsilon)n\Big].$$

If the value of M is sufficiently large depending on ε, this probability is larger than $1 - \exp(-n)$, as claimed. Proposition 4.24 for infinite second moment is proved. □

4.3.2 Nonzero mean

Now we will prove the part of Proposition 4.24 about nonzero mean. We can assume here that the second moment of the entries A_{ij} is finite, as the opposite case was treated in

Section 4.3.1. As before, we will first focus on one submatrix. In the following lemma we make an extra boundedness assumption, which we will get rid of using truncation later.

Lemma 4.26. *Consider an $m \times m$ random matrix B whose entries are i.i.d. random variables that satisfy*

$$\mathbb{E}B_{ij} = \mu > 0, \quad \mathbb{E}B_{ij}^2 \le \sigma^2, \quad |B_{ij}| \le K\sqrt{m} \text{ a.s.}$$

Then, for any $M > 0$ there exists m_0 that may depend only on μ, σ, K and M, and such that for any $m > m_0$ we have

$$\|B\| \ge \frac{\mu m}{2}$$

with probability at least $1 - \exp(-M^2 m)$.

Proof. Notice that

$$\|B\| \ge \frac{1}{m} \sum_{i,j=1}^{m} B_{ij}.$$

(To check this inequality, recall that $\|B\| \ge x^\mathsf{T} Bx$ for any unit vector x; use this for the vector x whose all coordinates equal $1/\sqrt{m}$.) Then we can bound the failure probability as follows:

$$\mathbb{P}\left\{\|B\| < \frac{\mu m}{2}\right\} \le \mathbb{P}\left\{\sum_{i,j=1}^{m} B_{ij} < \frac{\mu m^2}{2}\right\} \le \mathbb{P}\left\{\sum_{i,j=1}^{m} (B_{ij} - \mathbb{E}B_{ij}) < -\frac{\mu m^2}{2}\right\}$$

where we used that $\mathbb{E}B_{ij} = \mu$ in the last step.

Apply Bernstein's inequality for the random variables B_{ij} and use that they have variance at most σ^2 and are bounded by $K\sqrt{m}$ by assumption. The failure probability gets bounded by

$$\exp\Big(-\frac{\mu^2 m^4/8}{\sigma^2 m^2 + K\sqrt{m}/3}\Big).$$

If m is large enough depending μ, σ, K and M, then this probability can be further bounded by $\exp(-M^2 m)$, as claimed. □

Next, we will use truncation to get rid of the boundedness assumption in Lemma 4.26 and thus prove the following.

Lemma 4.27. *Consider an $m \times m$ random matrix B whose entries are i.i.d. random variables that satisfy*

$$\mathbb{E}B_{ij} = \mu > 0, \quad \mathbb{E}B_{ij}^2 \le \sigma^2.$$

Then, for any $M > 0$ there exists m_0 that may depend only on μ, σ, K, M and the distribution of the entries, and such that for any $m > m_0$ we have

$$\|B\| \geq M\sqrt{m} \tag{4.52}$$

with probability at least $1 - \exp(-M^2 m)$.

Proof. Choosing m_0 large enough depending on M and the distribution of B_{ij}, we can make sure that for any $m \geq m_0$ the truncated random variables

$$\bar{B}_{ij} := B_{ij}\mathbb{1}_{|B_{ij}|\leq M\sqrt{m}} \quad \text{satisfy} \quad \mathbb{E}\bar{B}_{ij} \geq \mathbb{E}B_{ij} - \frac{\mu}{2} = \frac{\mu}{2}.$$

(This follows easily from Lebesgue's monotone convergence theorem.)

Let us consider the event that all entries of B are appropriately bounded:

$$\mathcal{E} := \left\{|B_{ij}| \leq M\sqrt{m} \text{ for all } i, j \in [n]\right\}.$$

Suppose for a moment that (4.52) fails, so we have $\|B\| < M\sqrt{m}$. Since the inequality $\|B\| \geq \max_{i,j} |B_{ij}|$ is always true, the event $\mathcal{E}$ must hold in this case. This in turn implies that the truncation has no effect on the entries, i.e. $\bar{B}_{ij} = B_{ij}$ for all i, j.

We have shown that in the event of the failure of (4.52), we may automatically assume that the entries of B are appropriately bounded. Therefore the failure probability satisfies

$$\mathbb{P}\left\{\|B\| < M\sqrt{m}\right\} = \mathbb{P}\left\{\|\bar{B}\| < M\sqrt{m}\right\}$$

where $\bar{B}$ denotes the matrix with the truncated entries $\bar{B}_{ij}$. It remains to apply Lemma 4.26 for the random matrix $\bar{B}$, noting that truncation may only decrease the second moment. The failure probability gets bounded by $\exp(-M^2 m)$, as claimed. $\square$

Proof of Proposition 4.24 for non-zero mean. As we mentioned in the beginning of this section, we can assume that the entries B_{ij} have finite second moment σ^2. Then the conclusion of the proposition follows by exact same union bound argument as in the end of Section 4.3.1 (just use Lemma 4.27 instead of Lemma 4.25 there.) $\square$

4.3.3 Proof of Theorem 4.3

We will prove a stronger fact that

$$\min \frac{\|A'_n\|}{\sqrt{n}} \to \infty \quad \text{as} \quad n \to \infty \quad \text{almost surely,} \tag{4.53}$$

where the minimum is taken over all $(1-\varepsilon)n \times (1-\varepsilon)n$ submatrices A'_n of A_n. As we mentioned below Proposition 4.24, this would imply the conclusion of Theorem 4.3, since modifying an $\varepsilon n \times \varepsilon n$ submatrix leaves some $(1-\varepsilon)n \times (1-\varepsilon)n$ sub-matrix intact.

Fix any $M > 0$ and consider the events

$$\mathcal{E}_n := \Big\{ \min \frac{\|A'_n\|}{\sqrt{n}} \geq M \Big\}, \quad n = 1, 2, \ldots$$

where the minimum has the same meaning as before. By Proposition 4.24, there exists n_0 such that

$$\mathbb{P}(\mathcal{E}_n^c) \leq e^{-n} \text{ for all } n > n_0.$$

In particular, the series $\sum_{n=1}^{\infty} \mathbb{P}(\mathcal{E}_n^c)$ converges. The Borel-Cantelli lemma then implies that the probability that infinitely many $\mathcal{E}_n^c$ occur is 0. Equivalently, with probability 1 there exists N such that $\mathcal{E}_n$ hold for all $n \geq N$.

We have shown that for any $M > 0$, with probability 1 there exists N such that

$$\min \frac{\|A'_n\|}{\sqrt{n}} \geq M \quad \text{for all } n \geq N.$$

Intersecting these almost sure events for $M = 1, 2, \ldots$, we conclude (4.53). Theorem 4.3 is proved. □

4.4 Discussion of Theorem 4.1

There are several possible extensions of Theorem 4.1. Let us list some of them and then discuss the constructiveness question in more details.

1. It is natural to expect a version Theorem 4.1 even if the entries of A are *not identically distributed*. Our argument relies on the identical distribution in several places, including discretization arguments (proof of Theorem 4.6) and symmetrization (proofs of Lemmas 4.10 and 4.11).

2. A version of Theorem 4.1 should hold for *symmetric matrices* A with independent en-

tries on and above the diagonal. A simplest way to get this result would be to use Theorem 4.1 to control the parts of A above and below the diagonal separately, and then combine them. However, for this argument one would need a version of Theorem 4.1 for non-identical distributed entries.

3. It would be good to *remove the logarithmic factor* $\ln \varepsilon^{-1}$ from the bound in Theorem 4.1, or to show that this factor is necessary. Such bound would be optimal up to an absolute constant factor.

4. Finally, while Remark 4.2 states that *the dependence on* ε in Theorem 4.1 is optimal in general, this dependence might be dramatically improved under a natural boundedness assumption. Namely, suppose that the entries of A are $O(\sqrt{n})$ almost surely. (In fact, most of the proof – until Section 4.2.4 – was done under this additional assumption.) In this case, is the dependence of the norm on ε *logarithmic* in Theorem 4.1, i.e.

$$\|\tilde{A}\| \leq C \ln(\varepsilon^{-1})\sqrt{n}? \tag{4.54}$$

In fact, for the partial case of Bernoulli matrices such that $np = c_0 = const$ (where p is a probability of a non-zero entry) this bound can be quickly deduced from Corollary 4.18.

Indeed, after renormalization that imposes matrix elements to have variance one (so we deal with the scaled Bernoulli matrix with $B_{ij} = O(p^{-1/2})$), we can see that such matrices satisfy the boundedness assumption, as $B_{ij} = O(p^{-1/2}) = O(\sqrt{n}/\sqrt{c_0}) = O(\sqrt{n})$. Then, by Corollary 4.18 after a deletion of $\varepsilon n \times \varepsilon n$ submatrix we get a matrix $\tilde{B}$ with all rows and columns having at most

$$21pn + 4\ln \varepsilon^{-1} \leq 100 c_0 \ln \varepsilon^{-1}$$

non-zero elements of order $O(\sqrt{n})$. Hence,

$$\max_{i,j} \left(\|\tilde{B}_i\|_1,\, \|\tilde{B}^j\|_1 \right) \lesssim \sqrt{n} \cdot \ln \varepsilon^{-1}.$$

Applying Lemma 2.19 leads to (4.54).

5. Theorem 4.1 does not indicate what sub-matrix should be removed to improve the norm; it is rather an existential result. It would be nice to have an *explicit description of a submatrix to be removed.*

The last question was resolved in some partial cases in the previous related work. For example, for Bernoulli random variables, a variant of Theorem 4.1 was proved by U. Feige and E. Ofek [FO05]; see [LLV17] for an alternative argument and more general way to regularize such matrices. Suppose the entries of an $n \times n$ matrix B are independent Bernoulli random variables with mean $p \in (0, 1)$. If one removes the heavy rows and columns – those containing more than $2pn$ ones, then the resulting matrix B' satisfies the optimal norm bound $\|B' - \mathbb{E}B'\| = O(\sqrt{pn})$. To see that this bound is consistent with that of Theorem 4.1, divide both sides by $\sqrt{p}$ to normalize the variance of the entries. Moreover, one can quickly check using concentration that the number of heavy rows and columns in B is typically small. With a little more work, one can even place all ones from the heavy rows and columns into a small submatrix (see Lemma 4.17 and Section 5.2 below). Thus, Feige-Ofek's result is an example of Theorem 4.1, and in this example we actually have an explicit recipe of regularization: removal of the heavy rows and columns. Note, however, that the results of [FO05, LLV17] hold for symmetric matrices as well, while we do not know how to immediately extend Theorem 4.1 for symmetric matrices (the requirement of the identical distribution of entries of A prevents doing simple symmetrization tricks).

Weaker (constructive) versions of Theorem 4.1, with an additional factor $\log n$ in the norm bound and weaker probability guarantees, can be derived from known general bounds on random matrices, such as the matrix Bernstein's inequality (Lemma 2.8). One would apply the matrix Bernstein's inequality for the entries truncated at level $\sqrt{n}$, and control the larger entries as in Section 4.2.4.

In the next chapter we discuss several more approaches to a constructive "improvement" of the random matrix structure, in particular, to the regularization of its operator norm.

CHAPTER 5

Constructive regularization

One may naturally wonder what exactly causes the norm of a mean zero random matrix A to be too large. A natural guess is that the only troublemakers are a few large entries of A. Indeed, this is exactly how the necessity of the fourth moment for (4.1) was shown in [BSY88, Sil89]. So we may ask – can we obtain a result like Theorem 4.1 simply by zeroing out a few largest entries of A?

The answer is no. A counterexample is a sparse Bernoulli matrix A, whose i.i.d. entries take values $\pm\sqrt{n}$ with probability $1/2n$ each and 0 with probability $1-1/2n$. It is not hard to check that A is likely to have a row whose norm exceeds $c\sqrt{n}\log n/\log\log n \gg \sqrt{n}$, and consequently we have $\|A\| \gg \sqrt{n}$. In other words, without removal of any entries the norm of A is too large. However, if we are to remove any entries based purely on their magnitudes, we must remove them all. (Recall that all non-zero elements of A have the same magnitude $\sqrt{n}$.) But removal of all nonzero entries of A is not a local intervention, since such entries can not be placed in a small submatrix (as we explained in Remark 4.2).

In Chapter 5 we will discuss some algorithms that bound the operator norm of a random matrix, changing only a small fraction of the matrix entries. First, if the matrix entries have more than two bounded moments, then the approach discussed above actually works: we can zero out a few of the largest entries of the matrix A to make the norm $O(\sqrt{n})$. This is demonstrated in Section 5.1.

Then, if the entries of the matrix are scaled Bernoulli random variables (e.g. like in the counterexample above) a simple regularization procedure is possible: to ensure that the norm is $O(\sqrt{n})$, it is enough to zero out (or otherwise reweigh) a small fraction of the matrix rows and columns that have too many non-zero entries. This was proved earlier in [FO05, LLV17]. In Section 5.2 we improve upon these results by showing how to construct a small sub-matrix to zero out for the norm regularization. We also obtain a nice structural description of the sub-matrix to be deleted (see Theorem 5.3).

Finally, in Section 5.3 we state and prove the regularization algorithm that allows to work with general matrices which entries have only two finite moments. It allows to bound the norm to the order $\sqrt{\log\log n \cdot n}$ with high probability after zeroing out $n\varepsilon$ rows and columns of the matrix for any $\varepsilon > 0$ of our choice (see Theorem 5.14).

5.1 More than two finite moments

Under slightly stronger moment assumptions than in Theorem 4.1, zeroing out a few large entires does bring the norm of A down. The following result can be quickly deduced by truncation from known bounds on random matrices such as [vH17a, Seg00, AT16].

Theorem 5.1 ($2+\varepsilon$ moments)**.** *For any $\varepsilon \in (0,1]$ there exists $n_0 = n_0(\varepsilon)$ such that the following holds for any $n > n_0(\varepsilon)$. Consider an $n \times n$ random matrix A with i.i.d. mean zero entries which satisfy $\mathbb{E}|A_{ij}|^{2+\varepsilon} \leq 1$. Then, with the probability at least $1 - 2\exp(-n^{\varepsilon/5})$, there exists a integer $K \leq n^{1-\varepsilon/9}$ such that the matrix $\tilde{A}$ obtained by zeroing out K largest entries of A satisfies*

$$\|\tilde{A}\| \leq 9\sqrt{n}.$$

We will deduce Proposition 5.1 from the following general bound [BVH16, Remark 3.13] (version for rectangular matrices):

Theorem 5.2 (Bandeira-van Handel)**.** *Let X be an $n \times n$ matrix whose entries X_{ij} are independent centered random variables. Then there exists for any $\varepsilon \in (0, 1/2]$ a constant c_ε such that for every $t \geq 0$*

$$\mathbb{P}\{\|X\| \geq (1+\varepsilon)6\sigma + t\} \leq n\exp(-t^2/c_\varepsilon\sigma_*^2),$$

where

$$\sigma := \max(\sigma_1, \sigma_2), \quad \textit{where} \quad \sigma_1 = \max_i \sum_j \mathbb{E}(X_{ij}^2), \quad \sigma_2 = \max_j \sum_i \mathbb{E}(X_{ij}^2);$$

$$\sigma_* := \max_{ij} \|X_{ij}\|_\infty.$$

Proof of Theorem 5.1. Let us call an entry A_{ij} *large* if $|A_{ij}| > R := n^{1/2-\varepsilon/8}$, otherwise call the entry *small*.

We claim that there are very few large entries with high probability, and we can check this by the same argument as in Lemma 4.21. Indeed, the $2+\varepsilon$ moment assumption and

Chebyshev's inequality give

$$\mathbb{P}\left\{A_{ij} \text{ is large}\right\} = \mathbb{P}\left\{|A_{ij}| > R\right\} < \frac{1}{R^{2+\varepsilon}} \le \frac{1}{n^{1+\varepsilon/8}}, \tag{5.1}$$

where the last inequality follows by our choice of R. Thus the expected number of large entries is at most $n^2/n^{1+\varepsilon/8} = n^{1-\varepsilon/8}$. A standard application of Chernoff's inequality (see e.g. [Ver16, Chapter 2]) gives

$$\mathbb{P}\left\{A \text{ has more than } n^{1-\varepsilon/9} \text{ large entries}\right\} \le \left(\frac{en^{1-\varepsilon/8}}{n^{1-\varepsilon/9}}\right)^{n^{1-\varepsilon/9}},$$

which can be further bounded by $\exp(-n^{1/2})$ if n is sufficiently large in terms of ε. Hence we can zero out all large large entries of A. It remains to show that the result of this operation, which we denote by $\tilde{A}$, has norm at most $8\sqrt{n}$ with high probability.

For convenience, let us subtract the mean, and first bound

$$G := \tilde{A} - \mathbb{E}\tilde{A},$$

which is an $n \times n$ random matrix with independent mean zero entries G_{ij}. By Theorem 5.2, for any $t > 0$, we have

$$\mathbb{P}\{\|G\| \ge 7\sigma + t\} \le n\exp(-ct^2/\sigma_*^2), \tag{5.2}$$

where $c > 0$ is a constant and

$$\sigma^2 := \max\left\{\max_i \sum_j \mathbb{E}(G_{ij}^2), \max_j \sum_i \mathbb{E}(G_{ij}^2)\right\}, \qquad \sigma_* := \max_{ij} \|G_{ij}\|_\infty.$$

In our case,

$$\mathbb{E}G_{ij}^2 = \mathrm{Var}(\tilde{A}_{ij}) \le \mathbb{E}\tilde{A}_{ij}^2 \le \mathbb{E}A_{ij}^2 \le 1$$

(where we used the moment assumption) and

$$\|G_{ij}\|_\infty = \|\tilde{A}_{ij} - \mathbb{E}\tilde{A}_{ij}\|_\infty \le 2\|\tilde{A}_{ij}\|_\infty \le 2R.$$

Hence

$$\sigma \le \sqrt{n} \quad \text{and} \quad \sigma_* \le 2R.$$

Then, using (5.2) with $t = \sqrt{n}$, we conclude that

$$\mathbb{P}\{\|G\| \geq 8\sqrt{n}\} \leq n \exp(-cn/4R^2) \leq \exp(-n^{\varepsilon/5}) \tag{5.3}$$

where the last inequality holds due to the definition of R, if n is sufficiently large in terms of ε.

Finally, we need to bound the contribution of the mean $\mathbb{E}\tilde{A}$ which we subtracted in defining G. This can be done by the exactly same argument as we used in the proof of Theorem 4.1 in Section 4.2.4. We repeat it here for completeness. Note that all entries of $\mathbb{E}\tilde{A}$ are the same, thus

$$\begin{aligned}
\|\mathbb{E}\tilde{A}\| &= n\big|\mathbb{E}\tilde{A}_{ij}\big| = n\big|\mathbb{E}(A_{ij} - \tilde{A}_{ij})\big| \quad \text{(since } \mathbb{E}A_{ij} = 0\text{)} \\
&= n\big|\mathbb{E}A_{ij}\mathbb{1}_{|A_{ij}|>R}\big| \quad \text{(by definition of } \tilde{A}\text{)} \\
&\leq n\big(\mathbb{E}A_{ij}^2\big)^{1/2}\big(\mathbb{P}\left\{|A_{ij}| > R\right\}\big)^{1/2} \quad \text{(by Cauchy-Schwarz)} \\
&\leq n \cdot 1 \cdot n^{-1/2} \leq n^{1/2},
\end{aligned} \tag{5.4}$$

where we used the moment assumption and a weaker form of the bound (5.1).

Concluding, it follows from (5.3) and (5.4) that with probability at least $\exp(-n^{\varepsilon/5})$, we have

$$\|\tilde{A}\| \leq \|\tilde{A} - \mathbb{E}\tilde{A}\| + \|\mathbb{E}\tilde{A}\| \leq 9\sqrt{n}.$$

The proof of Proposition 5.1 is complete. □

5.2 Bernoulli random matrices and random graphs

Random matrices with the entries distributed as scaled Bernoulli distributed is an another case when local regularization of the norm can be made in a simple constructive way.

We can view any symmetric $n \times n$ Bernoulli random matrix B as an adjacency matrix of a random graph (on n vertices, so that the edge between i and j is present if and only if $B_{ij} \neq 0$). In this graph interpretation, to regularize the norm of the adjacency matrix means to obtain a version of the original graph, for which certain graph algorithms (such as spectral methods for community detection) are guaranteed to work with high probability. This is why for the graph applications it is especially important to be able to regularize operator norm of the matrix locally and constructively.

The main part of this section is Theorem 5.3. It justifies the work of the Adjacency Matrix Decomposition algorithm, that finds a small sub-matrix of the original matrix (i.e., a small sub-graph of the original random graph), which one needs to zero out for the norm regularization. Another result of Theorem 5.3 is a description of the "bad" sub-graph to be deleted. It turns out that, with high probability, one can direct the edges of this subgraph without cycles, such that every its vertex will have constant outcoming degree.

We start the next section with a short overview of the norm regularization question in the graph interpretation. Namely, we will discuss concentration on Erdős-Rényi graphs, its connection to the graph density and to the number of bounded moments of the standardly scaled adjacency matrix. Then we state Theorem 5.3 and Adjacency Matrix Decomposition algorithm. Then, in Section 5.2.3, we will discuss in more details the connection to the earlier results, such as [LLV17, FO05]. After that, we will prove Theorem 5.3 in the Section 5.2.4.

5.2.1 Concentration on random graphs

We consider random graphs generated from inhomogeneous Erdős-Rényi model $G(n, (p_{ij}))$, that is, every graph has exactly n vertices and the edges are formed independently with given probabilities p_{ij}. This is a generalization of the classical Erdős-Rényi model $G(n, p)$ where all edge probabilities are equal to p.

One of the crucial applications of such random graphs is that they can be successfully used to model large (non-random) networks. One can use this model to infer various network properties, such as the network radius, the node degree distribution, or community structure. These inference procedure will work only if a single realization of the network (represented by the random adjacency matrix A) is close to the population mean, or the true model, $\mathbb{E}A$ (which is not known). If A is close to $\mathbb{E}A$, then the observed A retains some good properties of $\mathbb{E}A$ (such as eigenstructure, necessary for the spectral methods, or approximate block structure, revealing the communities within the network). Then the algorithms that proved to work on $\mathbb{E}A$ will work almost as efficiently on A, the matrix that we actually observe. See a survey [LLV18b] for much more detailed discussion of this property, also called *concentration of the random graph* and its applications to the network analysis.

Concentration of the random graph depends on a crucial parameter of the graph G, namely its, *generalized maximal expected degree*

$$d := \max_{ij} np_{ij}.$$

Clearly, d measures expected sparsity of a random graph. It is known that if the graph is *dense* (d is at least of order $O(\log n)$), then the concentration usually takes place. Indeed, the adjacency matrix A of the graph G with $d \geq C\log(n)$ satisfies with high probability

$$\|A - \mathbb{E}A\| = O(\sqrt{d}) \tag{5.5}$$

(see [FO05, LR15]). The estimate (5.5) means concentration as soon as $\|\mathbb{E}A\| \gg \sqrt{d}$. Note that this is the case for the classical Erdős-Rényi model (if $d = np$, then $\|\mathbb{E}A\| = d$) and for all the cases when p_{ij} are of similar order (so that maximum in the definition of d gives correct information about the expected degree of a vertex of G).

When d is less than $O(\log n)$ (graph G is too *sparse*) the estimate (5.5) does not take place and there is no good concentration. In this case the natural question is whether and how can we modify the original graph to achieve the same concentration for the regularized version:

$$\|\tilde{A} - \mathbb{E}A\| = O(\sqrt{d}). \tag{5.6}$$

Finally, consider for simplicity the standard Erdős-Rényi model $G(n,p)$ and let us normalize the adjacency matrix to have second moment one (the same scale we consider throughout the majority of this text, including Theorems 4.1 and 4.3). Then $B_{ij} := A_{ij}/\sqrt{p}$ and $\mathbb{E}B_{ij}^2 = \mathbb{E}A_{ij}^2/p \sim 1$. Note that (5.6) is equivalent to the familiar bound

$$\|\tilde{B} - \mathbb{E}B\| = O(\sqrt{n}),$$

and sparse regime $d \sim const$ corresponds to the case when $p \sim n^{-1}$, and $\mathbb{E}B_{ij}^{2+\varepsilon} \sim n^{\varepsilon/2}$ unbounded for $n \to \infty$ for any $\varepsilon > 0$. So, this is the case when the matrix elements have exactly two finite moments.

Keeping this in mind, we will stick to the original (5.6) normalization until the end of Section 5.2, as it is standard for the random graphs literature and was used in predecessor results such as [LLV17, FO05].

5.2.2 Graph regularization theorem and Adjacency Matrix Decomposition algorithm

We show that for the regularization (5.6) it is enough to delete a small sub-matrix A_0 of the adjacency matrix A, and that A_0 itself has a reasonable structure. Let us give some notations and then state our main decomposition theorem for random (adjacency) matrices.

If $A = (A_{ij})_{i,j=1}^n$, then we denote

$$A \setminus A_0 := \begin{cases} 0, & \text{if } (i,j) \in A_0; \\ A_{ij}, & \text{otherwise.} \end{cases}$$

By *symmetric permutation* (or symmetric swap) of the rows and columns of A we mean a pair of permutations, exchanging i-th row with j-th row of A, and then i-th column with j-th column. For two subsets of rows (and columns) we denote it $\{i_1, \ldots, i_k\} \leftrightarrow \{j_1, \ldots, j_k\}$.

Theorem 5.3. *Let $G = G(n, (p_{ij}))$ be an inhomogeneous Erdős-Rényi graph and d denotes its maximal expected degree $d = \max_{ij} np_{ij}$, $d \geq 5$ and let $r \geq 1$. Then for any n large enough with probability at least*

$$1 - 6 \cdot (10ne^{-d})^{-r} \tag{5.7}$$

there exists a symmetric permutation $\pi(A)$ of the graph adjacency matrix A with the following structure:

(a) *A top left $s \times s$ submatrix of $A_0 \subset \pi(A)$ has no more than $40r$ ones in each column above its main diagonal (and, from symmetry, as well in every row to the left from main diagonal), and $s \leq 10ne^{-d}$;*

(b) *All the rows and columns of $\pi(A) \setminus A_0$ have at most $12d$ ones. In particular, the matrix $\pi(A)$ concentrates on the complement of A_0, namely,*

$$\|\pi(A) \setminus A_0 - \mathbb{E}A\| \leq Cr^{3/2}\sqrt{d}.$$

Remark 5.4. In the proof of Theorem 5.3 obtain the following probability estimate for the good event:

$$1 - 3\max\{(10ne^{-d})^{-r}, \exp(-ne^{-d}/4)\}.$$

As it is shown in the end of Section 5.2.4 (proof of Theorem 5.3), the polynomial term dominates if we can take any

$$n > 27r^2e^d. \tag{5.8}$$

Note that inequality (5.8) is always satisfied for n large enough, given that $d \ll \log n$. Indeed, this implies $27r^2e^d \ll e^{\log n} = n$. And for $d \gtrsim \log n$ regularization is not needed, since, as we mentioned above,

$$\|A - \mathbb{E}A\| = O(\sqrt{d}) \text{ as long as } d = \Omega(\log n).$$

Algorithm 1: Adjacency Matrix Decomposition

Input: A - symmetric $n \times n$ matrix with $a_{ij} \in \{0, 1\}$, d, r

Output: $\pi(A)$ with the structure described in Theorem 5.3 and Figure 1

1. Find $\{i_1, \dots, i_{\xi_1}\}$ — all the rows of A with more than $12d$ ones;
   make a swap $\{i_1, \dots, i_{\xi_1}\} \leftrightarrow \{1, \dots, \xi_1\}$
2. Find $\{i_1, \dots, i_{\xi_2}\}$ — all non-empty columns inside the block $\{1, \dots, \xi_1\} \times [n]$;
   make a swap $\{j_1, \dots, j_{\xi_2}\} \leftrightarrow \{\xi_1 + 1, \dots, \xi_1 + \xi_2\}$

Let $A_0 := A_{\{1,\dots,\xi_1+\xi_2\}\times\{1,\dots,\xi_1+\xi_2\}}$. We expect to have concentration on $A \setminus A_0$.

3. Set $A' := A_0$, $k = 1$.
4. While [A' is not empty and $k \leq \lfloor \log_2(10ne^{-d}) \rfloor$] repeat
   4a. Find $\{i_1, \dots, i_{m_1}\}$ — all the rows of A' with more than $40r$ ones;
       make a swap $\{i_1, \dots, i_{m_1}\} \leftrightarrow \{1, \dots, m_1\}$
   4b. $A' := \{1, \dots, m_1\} \times \{1, \dots, m_1\}$,
       $k := k + 1$

Table 5.1: Regularization algorithm

Finally, note that in the most interesting case when d is a large constant independent from n (A is very sparse and yet the size of A_0 is a very small fraction of n), $n_0 := 5r^2e^d$ is a constant and it is enough to take $n \geq n_0$.

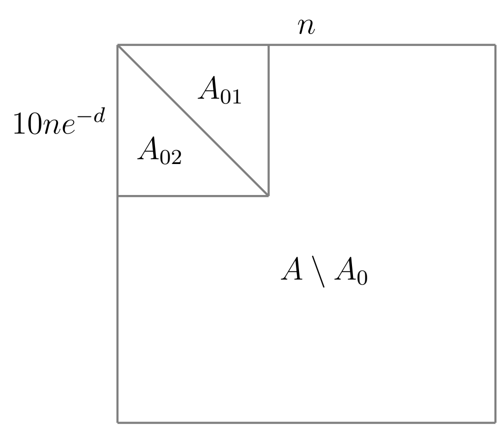

Figure 5.1: The adjacency matrix structure: short rows in A_{02}, short columns in A_{01}, concentration on the rest

Remark 5.5. A crucial feature of this theorem is that a symmetric permutation $\pi(A)$ with the claimed structure can be obtained in a constructive way for any sample graph $G(n, (p_{ij}))$, see Adjacency Matrix decomposition Algorithm below. If A is an adjacency matrix of a sample graph G taken form the probability model $G(n, (p_{ij}))$ with d being its maximal expected degree, then Theorem 5.3 justifies that the algorithm works with probability $1 - O((ne^{-d})^{-r})$.

Remark 5.6 (Graph interpretation). To conclude the first section, let us discuss what the result of Theorem 5.3 means in terms of vertices and edges of the sample graph $G = (V, E)$. Assume that G was taken from the inhomogeneous Erdős-Rényi model and d is its maximal expected degree.

Theorem 5.3 claims that with high probability $G = (V, E)$ can be decomposed into two subgraphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ in the following way.

Set of edges is decomposed into two non-intersecting subsets $E = E_1 \sqcup E_2$. Here, E_1 is a subset of the edges defined by A_0 being symmetric sub-matrix of the adjacency matrix of the graph (note that symmetric permutations of an adjacency matrix correspond to enumerations of the graph vertices, i.e. they do not change the graph itself). We define sets of vertices in a natural way:

$$V_1 = \{\text{all ends of the edges in } E_1\},$$

$$V_2 = \{\text{all ends of the edges in } E_2\}.$$

The subgraph $G_2 = (V_2, E_2)$ has all the vertices of degree at most $8d$, and its adjacency matrix satisfies concentration property (5.5). And the subgraph $G_1 = (V_1, E_1)$ is

- small: it has at most $10ne^{-d}$ vertices, all in 1-neighborhood of ”heavy” vertices, that is the vertices of G with the degree more than $12d$,
- its edges can be directed without cycles, such that every vertex in G_1 will have at most $40r$ outcoming edges.

Indeed, Theorem 5.3 provides us with the ordering of vertices of G_1 such that if we direct the edges E_1 according to this order, from smaller to larger index, every vertex will have at most $40r$ outcoming edges. It is enough to note that the collection of ways to direct the edges with respect to some enumeration coincides with the collection of ways to do it cycle free (see, for example, [Koz12]).

5.2.3 Prior work on random graphs regularization

Originally, it was proved in [FO05], that given a random symmetric Bernoulli matrix, zeroing out all its rows and columns with more than $O(d)$ non-zero elements brings the norm of the matrix to the order $O(\sqrt{d})$. Here, d is expected number of the non-zero elements in each row. In terms of random graphs, this means to delete all the edges adjacent to the high-degree vertices. However, if high-degree vertices are exactly the “important” vertices of the graph, cutting all the edges adjacent to them might change the structure of the graph too much.

The following more delicate regularization theorem was proved in [LLV17]:

Theorem 5.7 (Le-Levina-Vershynin)**.** *Consider a random graph from the inhomogeneous Erdős-Rényi model $G(n,(p_{ij}))$, and let $d = \max_{ij} np_{ij}$. For any $r \geq 1$, the following holds with probability at least $1 - n^{-r}$. Consider any subset consisting of at most $10n/d$ vertices, and reduce the weights of the edges incident to those vertices in an arbitrary way. Let d' be the maximal degree of the resulting graph. Then the adjacency matrix A' of the new (weighted) graph satisfies*

$$\|A' - \mathbb{E}A\| \leq Cr^{3/2}(\sqrt{d} + \sqrt{d'}).$$

Regularization procedure proposed by Theorem 5.7 is simple and flexible. Comparing to the regularization we propose in Theorem 5.3, the disadvantage of our approach is that we zero out a fraction of the entries, whereas Theorem 5.7 allows an arbitrary reweighing of the matrix entries. However, Theorem 5.3 shows that we can localize all the required changes to a small subgraph (which was not directly guaranteed by Theorem 5.7). Additional advantage of our result is the description of the "bad" submatrix A_0.

The idea of our work was inspired by the following decomposition theorem, used for the proof of Theorem 5.7 ([Theorem 2.6, [LLV17]]):

Lemma 5.8 (Adjacency Matrix Decomposition-1)**.** *Consider a random directed graph from taken the inhomogeneous Erdős-Rényi model, and let $d = \max_{ij} np_{ij}$. For any $r \geq 1$, with probability $1 - 3n^{-r}$, we can decompose the set of edges $[n] \times [n]$ into three classes N, R, C so that the following properties are satisfied for the adjacency matrix A:*

1. *The graph concentrates on N, namely, $\|(A - \mathbb{E}A)_N\| \leq Cr^{3/2}\sqrt{d}$.*
2. *Each row of A_R and each column of A_C has at most $32r$ ones.*

Moreover, all elements of R are within at most n/d columns of A, and all elements of C are within at most n/d rows of A.

In [LLV17] this theorem was used to identify the heaviest rows and columns of A (and then bound norms of A_R. A_C and A_H in the different ways). However, this result is of independent interest, as it describes the structure of a random graph, in particular, of its part containing the high-degree vertices. The disadvantage of Lemma 5.8 is that the decomposition was obtained in a completely non-constructive way: it finds a block N'

such that the norm $\|(A-\mathbb{E}A)_{N'}\|_{\infty\to 2}$ concentrates, and then uses Grothendieck-Pietsch factorization (see [Pie78, LT13] or [LLV17] for discussion) to conclude that the operator norm concentrates on a large sub-matrices of N'. The Grothendieck-Pietsch step makes it impossible to actually find classes N, R and C for any given sample matrix.

Our proof uses the known regularization result (Theorem 5.7) to improve the description of the set of heavy vertices. The observation playing a crucial role in this improvement is quite obvious by itself: with the high probability the quantity of heavy vertices of G is much smaller than n/d used in Theorem 5.7. The main step of the construction (an inductive decomposition procedure) is similar to the one used for the proof of Lemma 5.8.

Before we proceed to the proof, let us state an almost immediate corollary of the Theorem 5.7 (effectively a version of the Theorem 5.7 that is convenient for us for the later use):

Corollary 5.9. *Let $G = G(n,(p_{ij}))$ be an inhomogeneous Erdős-Rényi model and $d = \max_{ij} np_{ij}$. Then for any $r, d \geq 1$ with the probability at least $1-2n^{-r}$ the following holds. Let us consider all the vertices of the graph with degrees larger than $12d$ and reduce the weights of their vertices in arbitrary way such that all degrees become bounded by $12d$, then the adjacency matrix A' of a new graph concentrates:*

$$\|A' - \mathbb{E}A\| \leq Cr^{3/2}\sqrt{d}.$$

To derive this corollary from Theorem 5.7 it is easy to check that there are at most n/d rows having too many (more than $12d$) non-zero elements with probability at least $1-n^{-r}$. We will use only a simple version of Corollary 5.9, when the reduction of weights is performed by deletion of all edges adjacent to vertices with degrees more than $O(d)$.

5.2.4 Proof of the Theorem 5.3

The proof of Theorem 5.3 is constructive: it follows the steps of the Adjacency Matrix Decomposition algorithm (Table 5.1). First, we will prove three lemmas, estimating "probability losses" at Step 1, Step 2 and at one iteration of the Step 4 (see also Figure 2), and then combine their results in the end of the section.

Lemma 5.10 shows that the number of high-degree vertices in $G = G(n, p_{ij})$ is at most $O(ne^{-d}/d)$ with high probability (compare with $n/d \gg ne^{-d}/d$ that was used in [LLV17]):

Lemma 5.10. *Let A be an adjacency matrix of the graph $G(n, p_{ij})$ with the maximal expected degree d (i.e. $d = \max np_{ij}$), assume that $d \geq 5$. Then with the probability*

$$1 - 2\exp(-ne^{-d})$$

the number of the rows of A with more than $12d$ non-zero elements is at most ne^{-d}/d.

Proof. By definition, A is a symmetric matrix such that all its entries below main diagonal are independent Bernoulli random variables. Then if we define $A_1 := (\tilde{a}_{ij})_{i,j=1}^n$ and $A_2 := (\bar{a}_{ij})_{i,j=1}^n$ by

$$\tilde{a}_{ij} = \begin{cases} 0, & \text{if } i < j; \\ a_{ij}, & \text{otherwise;} \end{cases}$$

and

$$\bar{a}_{ij} = \begin{cases} 0, & \text{if } i > j; \\ a_{ij}, & \text{otherwise;} \end{cases}$$

then both A_1 and A_2 have independent entries and expected number of ones in any row of A_1 or A_2 is at most d, as

$$\sum_j p_{ij} n \leq \max_j p_{ij} \cdot n = d.$$

Hence, by Chernoff's inequality,

$$\mathbb{P}\{ \text{ a row of } A_1 \text{ has more than } 6d \text{ ones } \} \leq \left(\frac{ed}{6d}\right)^{6d} \leq e^{-4d}.$$

Let χ_1 be a set of rows of A_1 with more than $6d$ ones, $\chi_1 \subset \{1, \ldots, n\}$. Then

$$\mathbb{E}\{|\chi_1|\} = \mathbb{E}\sum_{i=1}^n \mathbb{1}_{\{ \text{ } i\text{-th row is in } \chi_1 \text{ } \}} = n \cdot \mathbb{P}\{ \text{ } i\text{-th row is in } \chi \text{ } \} \leq ne^{-4d}.$$

Again, by Chernoff's inequality,

$$\mathbb{P}\{|\chi_1| > \frac{ne^{-d}}{2d}\} \leq (\frac{2ende^{-4d}}{ne^{-d}})^{ne^{-d}/2d} \leq \exp(-ne^{-d}).$$

Similarly, if χ_2 is a subset of rows A_2, containing more than $6d$ ones, then

$$\mathbb{P}\{|\chi_2| > \frac{ne^{-d}}{2d}\} \leq \exp(-ne^{-d}).$$

Hence, with probability $1 - 2\exp(-ne^{-d})$ all but ne^{-d}/d rows of $A = A_1 + A_2$ contain at most $12d$ ones. $\square$

As the conclusion of this lemma and Corollary 5.9, we expect to have concentration on $n(1-e^{-d}/d)\times n(1-e^{-d}/d)$ sub-matrix A'. The next lemma shows that $A\setminus A'$ is very sparse (regardless of the fact that it contains the heaviest rows and columns).

Lemma 5.11. *Let A be an adjacency matrix of the graph $G(n,p_{ij})$ with the maximal expected degree $d=\max np_{ij}$. Let $d\geq 5$ and $m=ne^{-d}/d$. Then with the probability*

$$1-\exp(-ne^{-d}/4)$$

the following holds for all $m\times n$ blocks of A. If $B=I\times[n]$ is any such block ($I\subset\{1,\dots,n\}$, $|I|=m$), then its $m\times(n-m)$ part $B\setminus(I\times I)$ has at most $9ne^{-d}$ non-empty columns.

Proof. Consider one $m\times n$ block $B=I\times[n]\subset A$. Let $\bar{B}=B\setminus(I\times I)$, then $\bar{B}$ has independent entries and expected number of ones in every column of $\bar{B}$ is at most $e^{-d}<1$. Hence, by Chernoff's inequality,

$$\mathbb{P}\{\text{ a column in the block is non-empty}\}\leq\frac{e}{e^d}\leq e^{1-d}.$$

Let $\chi_{\bar{B}}$ be a set of non-empty columns in $\bar{B}$, $\chi_B\subset\{1,\dots,n\}\setminus I$ (by definition of $\bar{B}$ all the columns in I are empty). Then

$$\mathbb{E}\{|\chi_{\bar{B}}|\}=n\cdot\mathbb{P}\{\text{ }i\text{-th row is in }\chi_{\bar{B}}\text{ }\}\leq n\cdot e^{1-d}.$$

Again, by Chernoff's inequality and taking a union bound over all $m\times n$ blocks in A,

$$\begin{aligned}
&\mathbb{P}\{\text{ there exists a }m\times n\text{ block }A'\text{ with }|\chi_{A'}|>9ne^{-d}\}\\
&\qquad\leq\binom{n}{m}\cdot\mathbb{P}\{|\chi_{A'}|>9ne^{-d}\}\\
&\qquad\leq\left(\frac{n\cdot ed}{ne^{-d}}\right)^{ne^{-d}/d}\cdot\left(\frac{e^2}{9}\right)^{9ne^{-d}}\\
&\qquad\leq\exp\left[ne^{-d}\left(\frac{d+1}{d}+\frac{\ln d}{d}+18-9\ln 9\right)\right]\\
&\qquad\leq\exp(-ne^{-d}/4),
\end{aligned}$$

for all $d\geq 5$. Lemma 5.11 is proved.

$\square$

The next lemma uses the same technique as Lemmas 5.10 and 5.11. It is needed for the inductive decomposition of the non-empty part of heavy columns into two parts (the one with light rows and the one with light columns, see also Figure 3 below). It shows that with high probability any small square sub-block of A has constant number of ones in half of its columns.

Lemma 5.12. *Let A be an adjacency matrix of the graph $G(n, p_{ij})$ with the maximal expected degree $d = \max np_{ij}$, assume that $d \geq 5$. Let $k \in \mathbb{Z}_+$ and we consider square sub-matrices of the size $2m_k \times 2m_k$, where $m_k = \lfloor 10ne^{-d}2^{-k} \rfloor$. Let $r \geq 1$. Then with the probability*

$$1 - 2 \cdot \exp(-40 \ln 2 \cdot nrke^{-d}2^{-k})$$

all $2m_k \times 2m_k$ blocks of A have no more than m_k columns with at least $40r$ non-zero elements.

Proof. First, let us decompose A into the sum of two matrices (upper and lower triangular) with independent entries, as we did in Lemma 5.10, $A = A_1 + A_2$. Let us fix one $2m_k \times 2m_k$ block A' and estimate the expected number of heavy columns (i.e. containing at least $40r$ non-zero elements) in $A' \cap A_1$ and in $A' \cap A_2$ separately.

Consider a fixed $2m_k \times 2m_k$ sub-matrix. In both cases, the expected number of ones in one column is at most $2m_k p \leq 20de^{-d}2^{-k}$ and probability that some column contains more than $20r$ ones is bounded by Chernoff's inequality by

$$\left(\frac{e \cdot 20de^{-d}2^{-k}}{20r} \right)^{20r} \leq d^{20r} e^{(1-d)20r} 2^{-20kr}.$$

Let $\chi_{A_1'}$ be the set of columns in $A' \cap A_1$ with more than $20r$ ones, $\chi_{A_1'} \subset \{1, \ldots, 2m_k\}$. Then

$$\mathbb{E}\{|\chi_{A_1'}|\} \leq 20ne^{-d}2^{-k} \cdot d^{20r} e^{(1-d)20r} 2^{-20kr}.$$

By Chernoff's inequality again,

$$\mathbb{P}\{|\chi_{A_1'}| > m_k/2\} \leq 1 \cdot (4e \cdot e^{(1-d)20r} 2^{-20kr} d^{20r})^{5ne^{-d}2^{-k}}.$$

Clearly, the same estimate holds for the heavy rows of $A' \cap A_2$, so, if $\chi_{A'}$ be a set of columns in A' with more than $40r$ ones, then

$$\mathbb{P}\{|\chi_{A'}| > m_k\} \leq 2(4e \cdot e^{(1-d)20r} 2^{-20kr} d^{20r})^{5ne^{-d}2^{-k}}$$

Now, taking a union bound over all $2m_k \times 2m_k$ blocks in A,

$$\begin{aligned}
&\mathbb{P}\{\text{ there exists a } 2m_k \times 2m_k \text{ block } A' \text{ with } |\chi_{A'}| > 10ne^{-d}2^{-k}\} \\
&\le \binom{n}{2m_k}^2 \cdot \mathbb{P}\{|\chi_{A'}| > 10ne^{-d}2^{-k}\} \\
&\le 2\left[\left(\frac{n \cdot e}{20ne^{-d}2^{-k}}\right)^8 \cdot 4e \cdot e^{(1-d)20r}2^{-20kr}d^{20r}\right]^{5ne^{-d}2^{-k}} \\
&\le 2\left[2^{k(8-16r)} \cdot 2^{-4kr}\frac{4e}{20^8}e^{(1+d)8+(1-d)20r}d^{20r}\right]^{5ne^{-d}2^{-k}}. \qquad (5.9)
\end{aligned}$$

We will use that

$$2^{-4kr}\frac{4e}{20^8}e^{(1+d)8+(1-d)20r}d^{20r} \le 1$$

for any $k \ge 1$, $r \ge 1$ and $d \ge 5$. Indeed, $\frac{4e^8}{20^8} < 1$ and to check

$$\exp(-4\ln 2kr + 8d + 20r - 20dr + 20r\ln d) \le 1$$

it is enough to show that

$$8d + 18r - 20dr + 20r\ln d \le 0 \qquad (5.10)$$

for r and d as considered (since $-4\ln 2kr \le -2r$ for $k \ge 1$). Finally, (5.10) is satisfied, as for $d \ge 5$

$$\begin{aligned}
8d + 18r - 20dr + 20r\ln d &\le 8d + 18r - 20r(d - d/3)) \\
&= 8d + 18r - 40/3dr \\
&\le (8 - 40/3)dr + 18r \\
&\le (18 - 16/3d)r < 0, \qquad (5.11)
\end{aligned}$$

as the product of a negative and a positive numbers.

Hence,

$$(5.9) \le 2\left[2^{k(8-16r)}\right]^{5ne^{-d}2^{-k}} \le 2\exp[-40\ln 2rkne^{-d}2^{-k}].$$

This concludes the proof of Lemma 5.12. □

Now we are ready to combine all probability estimates obtained to prove that Adjacency Matrix Decomposition algorithm achieves the claimed results, i.e. to prove Theorem 5.3.

Proof of Theorem 5.3. First, let us find all the rows of A with more than $12d$ ones and perform a symmetric swap, placing then on the top of the matrix. Let $\{1,\dots,\xi_1\}$ be their new indices. Then, consider columns of A that are non-empty within the block $\{1,\dots,\xi_1\}\times\{\xi_1+1,\dots,n\}$. Let us place them into positions $\{\xi_1+1,\dots,\xi_2\}$ by another symmetric swap. We define $A_0 := \{1,\dots,\xi_2\}\times\{1,\dots,\xi_2\}$.

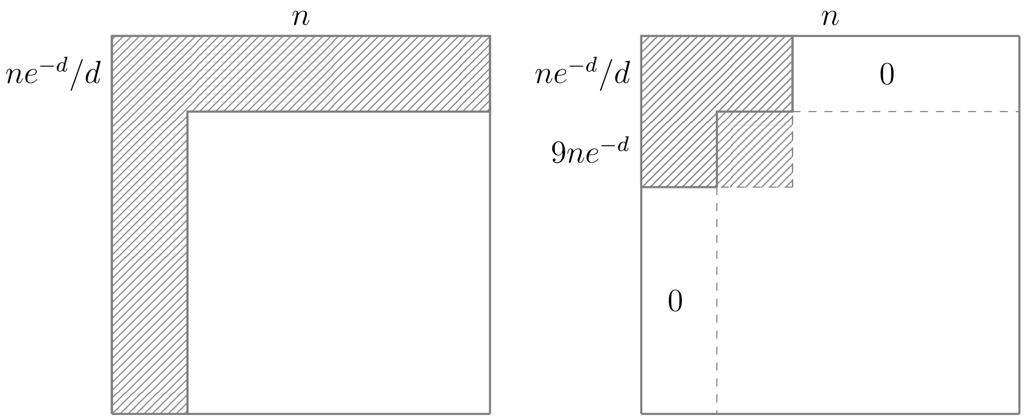

Step 1: heavy rows/columns are dashed   Step 2: A_0 is dashed

Figure 5.2: Construction of A_0

So, all the rows and columns with the indices $\{\xi_2+1,\dots,n\}$ contain at most $10d$ ones by construction, and this property will be preserved as all further symmetric permutations will happen within rows and columns with the indices in $\{1,\dots,\xi_2\}$. Hence, by Corollary 5.9 we will have the concentration outside A_0:

$$\mathbb{P}\{\|\pi(A)\setminus A_0-\mathbb{E}A\|>Cr^{3/2}\sqrt{d}\}\le 2n^{-r}, \tag{5.12}$$

so, condition (b) of Theorem 5.3 is satisfied.

Also note that with high probability the size of A_0 is at most $m\times m$, where $m=10ne^{-d}$. Indeed, consider exceptional events

$$\tilde{\mathcal{E}}_1 := \{\xi_1 > ne^{-d}/d\},\qquad \tilde{\mathcal{E}}_2 := \{\xi_2-\xi_1 > 9ne^{-d}\}.$$

By Lemmas 5.10 and 5.11

$$\mathbb{P}(\tilde{\mathcal{E}}_1)\le 2\exp(-ne^{-d}),\qquad \mathbb{P}_{\tilde{\mathcal{E}}_1}(\tilde{\mathcal{E}}_2)\le \exp(-ne^{-d}/4),$$

and

$$\mathbb{P}(\xi_2>10ne^{-d})=\mathbb{P}(\tilde{\mathcal{E}}_1\cup\tilde{\mathcal{E}}_2)\le 3\exp(-ne^{-d}/4). \tag{5.13}$$

To conclude the proof of Theorem 5.3 we just need to permute the rows and columns of A_0 to satisfy the structure condition (a). We proceed with the algorithm: set $A' := A_0$, find all the columns in A' with more than $40r$ ones, make a symmetric swap assigning them indices $\{1, \dots, m_1\}$ (for some $m_1 \leq m$). Then note that due to the symmetry of A', the restriction $A' \setminus (\{1, \dots, m_1\} \times \{1, \dots, m_1\})$ already satisfies the structure condition (a). So we redefine A' to be just the exceptional minor, $A' := \{1, \dots, m_1\} \times \{1, \dots, m_1\}$ and repeat with the new A' (see Figure 3 below).

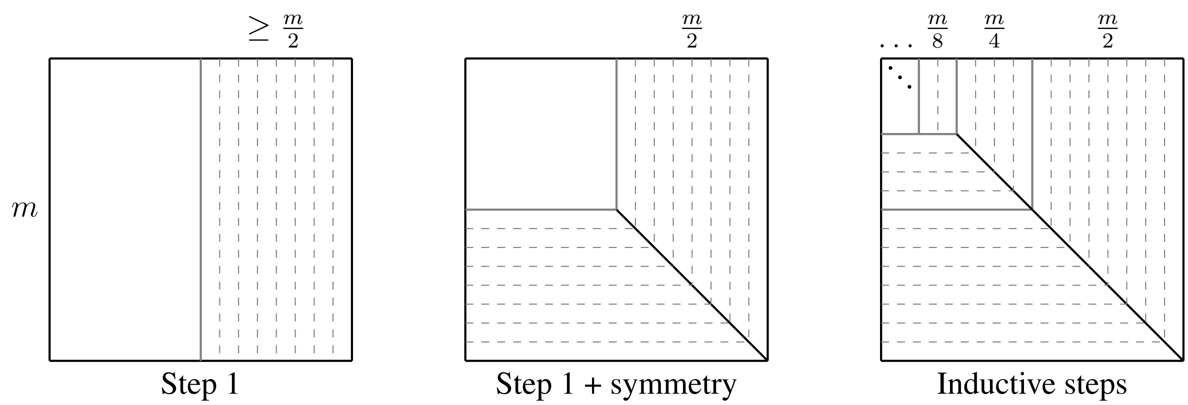

Figure 5.3: Inductive decomposition A_0. Dotted lines – "light" rows or columns

By definition, the algorithm makes at most $k = \lfloor \log_2(10ne^{-d}) \rfloor$ steps and, clearly, in the end condition (a) is satisfied for the matrix $A_0 \setminus A'$. So, we just need to show that with high probability A' will be empty after k steps.

To estimate this probability, note that if the size of the matrix A' decreases at least twice at every step, then on the complement of $\tilde{\mathcal{E}}_1 \cup \tilde{\mathcal{E}}_2$ we have as required: $A' = 0$ after k steps. Let us introduce several notations:

$$A'_j := A' \text{ after } j \text{ steps of the algorithm.}$$

So, we obtain a $m_j \times m_j$ matrix A'_j for $j = 1, \dots, k$. Let $\mathcal{E}_j$ be the j-th exceptional event:

$$\mathcal{E}_j := \{ \text{ there are } m_j/2 \text{ columns with more that } 40r \text{ ones in } A_j \ \}.$$

Then we need to bound $\mathbb{P}(\mathcal{E}) = \mathbb{P}(\bigcup_{j=0}^{k} \mathcal{E}_j)$. Note that on $\bigcap_{i=0}^{j-1} \mathcal{E}_i^c$ the size of A_j is bounded above $m_j \leq 10ne^{-d}2^{-j}$. So, by Lemma 5.12,

$$\mathbb{P}(\mathcal{E}_j \cap \bigcap_{i=1}^{j-1} \mathcal{E}_i^c) \leq 2\exp(-40 \ln 2 \cdot nrje^{-d}2^{-j}).$$

Therefore, we can estimate the probability of the exceptional event

$$\mathcal{E} := \bigcup_{j=1}^{k} \mathcal{E}_j = \bigcup_{j=1}^{k} \left(\mathcal{E}_j \cap \bigcap_{i=1}^{j-1} \mathcal{E}_i^c \right)$$

using the union bound:

$$\mathbb{P}(\mathcal{E}) \leq 2 \sum_{j=1}^{k} \exp(-40 \ln 2 \cdot rjne^{-d} 2^{-j}).$$

Then, let us upper bound every member of the sum by the maximal one, which is the one with $j = k$, since $j \cdot 2^{-j}$ is a decreasing function for $j \in Z_+$. Recall that $k = \lfloor \log_2(10ne^{-d}) \rfloor$. So,

$$\begin{aligned}
\mathbb{P}(\mathcal{E}) \leq 2 \log_2(10ne^{-d}) \cdot \exp(-40 \ln 2 \cdot rne^{-d} \log_2(10ne^{-d})(10ne^{-d})^{-1}) \\
\leq 2 \log_2(10ne^{-d}) \cdot \exp(-4 \ln 2 \cdot r \log_2(10ne^{-d})) \\
\leq 2 \frac{\log_2(10ne^{-d})}{(10ne^{-d})^{2r}} \\
\leq \frac{2}{(10ne^{-d})^r}.
\end{aligned} \tag{5.14}$$

So, by (5.12), (5.13) and (5.14) we will get to an adjacency matrix satisfying both conditions (a) and (b) after $\lfloor \log(10ne^{-d}) \rfloor$ steps of the Algorithm 1 with the probability at least

$$1 - 2n^{-r} - 3 \exp(-ne^{-d}/4) - 2(10ne^{-d})^{-r} \geq 1 - 6 \max\{(10ne^{-d})^{-r}, \exp(-ne^{-d}/4)\}.$$

The polynomial term dominates as long as n is large enough, in particular, $n > 27r^2 e^d$. This follows from the fact that

$$(10m)^r \leq \exp(m/4) \tag{5.15}$$

when $m = 27r^2$, and that the right hand side grows faster with m for $m > 27r^2$. To check (5.15), let us note that $x > 8 \ln x$ for $x \geq 27$. As $r \geq 1$, we have that $8 \ln(27r) \leq 27r$, thus, $8r \ln(\sqrt{270}r) \leq 27r^2$. Finally, this implies

$$r \ln 10 \cdot 27r^2 \leq \frac{27r^2}{4},$$

which is exactly (5.15) with $m = 27r^2$. This concludes the proof of Theorem 5.3. $\square$

5.3 General matrices with exactly two finite moments

Now, suppose that we have a general matrix with the entries having zero mean and finite second moment (no specific distribution assumptions, or extra moments assumptions are made). Theorem 4.1 does not provide an answer how to find the small submatrix to be deleted – it is rather an existential result. So, what can we do to regularize the norm?

As it was already mentioned in the end of Section 4.4, it is easy to show that deletion of the large entries regularizes the norm to the order $\sqrt{n} \cdot \ln n$. Indeed,

Lemma 5.13. *Consider an $n \times n$ random matrix A with i.i.d. entries that have zero mean and unit variance, and let $\varepsilon \in (0, 1]$. Then, for any $c \geq 1$, with probability at least $1-4n^{1-c}$, one can replace with zeros an $\varepsilon n \times \varepsilon n$ submatrix of A containing all elements A_{ij} such that $|A_{ij}| > \frac{5\sqrt{n}}{\sqrt{\varepsilon}}$ to get a matrix $\tilde{A}$ with the norm*

$$\|\tilde{A}\| \leq 10c\frac{\sqrt{n}}{\sqrt{\varepsilon}} \cdot \ln n.$$

Proof. By Lemma 4.21, with probability $1 - \exp(-\varepsilon n)$, all but εn entries of the matrix A are large such that $|A_{ij}| \leq \frac{5\sqrt{n}}{\sqrt{\varepsilon}}$. Let us zero out an $\varepsilon n \times \varepsilon n$ submatrix containing all these large entries. Let $\tilde{A}$ be the resulting matrix.

Exactly in the same way as we have checked in the proof Theorem 4.1 in Section 4.2.4, $\|\mathbb{E}\tilde{A}\| \leq \sqrt{n}$ and by triangle inequality

$$\|\tilde{A}\| \leq \|B\| + \|\mathbb{E}\tilde{A}\|, \quad \text{where} \quad B := \tilde{A} - \mathbb{E}\tilde{A}.$$

The matrix B has zero mean elements such that $|B_{ij}| \leq \frac{10\sqrt{n}}{\sqrt{\varepsilon}}$ and $\mathbb{E}B_{ij}^2 \leq 1$ (also like in the proof of Theorem 4.1). Note that

$$B = \sum_{i,j} E^{ij},$$

where E^{ij} are an $n \times n$ matrices with one non-zero element each:

$$E^{ij} := (e_{kl}^{ij})_{k,l=1}^n = \begin{cases} B_{kl}, & \text{if } i = k, j = l \\ 0, & \text{otherwise.} \end{cases}$$

Let us apply matrix Bernstein concentration inequality (Corollary 2.9) to the $\sum_{ij} E^{ij}$. For

$r \geq 1$ (to be specified below), take $t = r\sqrt{n}\ln n$, then

$$\mathbb{P}\left\{\left\|\sum_{i,j} E^{ij}\right\| \geq r\sqrt{n}\ln n\right\} \leq 4n\exp\left(-\frac{r^2 n \ln^2 n}{3\sigma^2 + \frac{2}{3}Kr\sqrt{n}\ln n}\right),$$

where $K = \max \|E^{ij}\|$ and $\sigma^2 = \max\left(\left\|\sum_{ij}\mathbb{E}E^{ij}(E^{ij})^T\right\|, \left\|\sum_{ij}\mathbb{E}(E^{ij})^T E^{ij}\right\|\right)$.

We have $\sigma^2 \leq n$ as the elements of the matrices $\mathbb{E}E^{ij}(E^{ij})^T$ and $\mathbb{E}(E^{ij})^T E^{ij}$ are $\mathbb{E}B_{ij}^2 \in [0,1]$. Also, $K = \|E^{ij}\| \leq \frac{10\sqrt{n}}{\sqrt{\varepsilon}}$, thus

$$\mathbb{P}\{\|B\| \geq r\sqrt{n}\ln n\} \leq 4n\exp(-\frac{r\sqrt{\varepsilon}}{10}\ln n).$$

Taking $r = 10c\varepsilon^{-1/2}$, we obtain

$$P\left\{\|B\| \geq 10c\frac{\sqrt{n}}{\sqrt{\varepsilon}}\ln n\right\} \leq 4n^{1-c}.$$

This concludes the proof of Lemma 5.13. $\square$

We can also relax extra $\ln n$ to extra $\sqrt{\ln\ln n}$ if we adopt more complex (but still local and constructive) regularization procedure.

Theorem 5.14 (Constructive regularization). *Suppose A is a random $n\times n$ matrix with i.i.d. symmetric entries A_{ij} such that $\mathbb{E}A_{ij}^2 = 1$. Let us denote 2^{-k}-quantiles of the distribution of $|A_{ij}|$ as*

$$q_k := \inf\{t : \mathbb{P}\{|A_{ij}| > t\} \leq 2^{-k}\}. \tag{5.16}$$

Let $\varepsilon \in (0, 1/2)$ and $\tilde{A}$ be the original matrix A after deletion of

- *elements A_{ij} such that $|A_{ij}| > q_{k_1}$*
- *rows and columns of A that contain more than $(c_\varepsilon 2^{-k} n)$ elements such that $q_{k-1} < |A_{ij}| \leq q_k$ for some $k \leq k_1$*

where $k_1 = \lceil \log_2 \frac{25n}{\varepsilon} \rceil$, $c_\varepsilon = 50/\varepsilon$. Then with probability $1 - 9n^{-9}$ the regularization described above is local (changes at most $2\varepsilon n$ rows $2\varepsilon n$ columns), and

$$\|\tilde{A}\| \leq Cc_\varepsilon\sqrt{n\cdot\ln\ln n}.$$

Here $C > 0$ is an absolute constant.

Theorem 5.14 will be proved in the next Section 5.3.1, followed by the discussion of the Theorem 5.14 conditions and potential improvements in the Section 5.3.2.

5.3.1 Proof of Theorem 5.14

General idea of our proof is the following. First, with high probability the regularization procedure described does not change the entries A_{ij} with $|A_{ij}| \ll \sqrt{n/\ln n} \sim q_{k_0}$ (Proposition 5.15). So, the matrix $S = \tilde{A} \cdot \mathbb{1}_{\{|A_{ij}| \leq q_{k_0}\}}$ has independent bounded entries and its norm is $O(\sqrt{n})$. This is shown in Lemma 5.16 (with a simple application of Theorem 5.2).

Because we also zero out the large entries such that $|A_{ij}| \geq q_{k_1}$, we are left to bound $M = \tilde{A} \cdot \mathbb{1}_{\{q_{k_0} < |A_{ij}| \leq q_{k_1}\}}$. For this part we use the proof strategy of Feige and Ofek (made for Bernoulli random matrices, see [FO05, CRV15]). Entries at every "level" $|A_{ij}| \in (q_{k-1}, q_k]$ are bounded by the entries of q_k-scaled Bernoulli matrices. We bound the norm of each of these Bernoulli matrices separately. Extra $\sqrt{\ln \ln n}$ factor appears when we sum up $\ln \ln n$ levels.

We will need several auxiliary results for the proof of Theorem 5.14. The first one checks that the regularization is likely to be local, and also does not change the entries in A that are small enough:

Proposition 5.15. *Let $\varepsilon \in (0, 1/2)$. Suppose A is a random $n \times n$ matrix with i.i.d. entries A_{ij} such that $\mathbb{E}A_{ij} = 0$, $\mathbb{E}A_{ij}^2 = 1$. Let q_k be the quantiles of the distribution of A_{ij} defined by* (5.16) *and*

$$e_{ki}^{row} := |j : |A_{ij}| \in (q_{k-1}, q_k]| \quad \textit{and} \quad e_{ki}^{col} := |j : |A_{ji}| \in (q_{k-1}, q_k]|$$

be the number of entries on the level k in i-th row or column of A. Then, let

$$R_{heavy}^K := \left|\{i : \exists k \in K, \textit{ such that } e_{ki}^{row} > c_\varepsilon \cdot 2^{-k} n\}\right|;$$

$$T_{heavy}^K := \left|\{i : \exists k \in K, \textit{ such that } e_{ki}^{col} > c_\varepsilon \cdot 2^{-k} n\}\right|$$

denote the number of "heavy" rows and columns on the levels $k \in K$. Let $k_1 = \lceil \log_2 \frac{25n}{\varepsilon} \rceil$ and $k_0 = \lfloor \log_2 \frac{c_1 n}{\ln n} \rfloor$ for some $c_1 < 1/12$. With probability at least $1 - 6n^{-10}$,

- $R_{heavy}^K < \varepsilon n$, $T_{heavy}^K < \varepsilon n$ *for* $K = \{k_0 + 1, \ldots, k_1\}$,
- $R_{heavy}^K = 0$, $T_{heavy}^K = 0$ *for* $K = \{1, \ldots, k_0\}$,
- $|A_{ij} : |A_{ij}| > q_{k_1}| \leq \varepsilon n$.

Proof. Fix a row i and a level $k \in \{k_0+1, \ldots, k_1\}$. By Chernoff's inequality (Lemma 2.10),

$$\mathbb{P}\{e_{ki}^{row} > c_\varepsilon \cdot 2^{-k} n\} \leq \left(\frac{e}{c_\varepsilon}\right)^{c_\varepsilon 2^{-k} n}.$$

Taking union bound over $k = k_0, \dots, k_1$,

$$\mathbb{P}\{\exists k \in \{k_0+1, \dots, k_1\} : e_{ki}^{row} > c_\varepsilon \cdot 2^{-k} n\} \le \sum_{k=k_0+1}^{k_1} \left(\frac{e}{c_\varepsilon}\right)^{c_\varepsilon 2^{-k} n}$$

$$\le 2 \cdot \left(\frac{e}{c_\varepsilon}\right)^{c_\varepsilon 2^{-k_1} n} \le 2 \cdot \left(\frac{e}{c_\varepsilon}\right)^{c_\varepsilon \varepsilon n/50n} \le 2 \cdot \frac{e\varepsilon}{50} < \frac{\varepsilon}{4},$$

if we take $c_\varepsilon = 50/\varepsilon$ (we used that geometric series in k can be it estimates by its largest term). Then,

$$\mathbb{E}|\{i : \exists k \in \{k_0+1, \dots, k_1\}, \text{ such that } e_{ki}^{row} > c_\varepsilon \cdot 2^{-k} n\}| \le \frac{n\varepsilon}{4}.$$

Finally, by Chernoff's inequality, for $K = \{k_0+1, \dots, k_1\}$,

$$\mathbb{P}\{|R_{heavy}^K| \ge n\varepsilon\} \le (e/4)^{n\varepsilon} \le \exp(-n\varepsilon/2).$$

Clearly the same estimate is true for the heavy columns.

Now, for $K = \{0, \dots, k_0\}$ we can repeat almost the same argument:

$$\mathbb{P}\{\exists k \in \{0, \dots, k_0\} : e_{ki}^{row} > c_\varepsilon \cdot 2^{-k} n\} \le \sum_{k=0}^{k_0} \left(\frac{e}{c_\varepsilon}\right)^{c_\varepsilon 2^{-k} n}$$

$$\le 2 \cdot \left(\frac{e}{c_\varepsilon}\right)^{c_\varepsilon 2^{-k_0} n} \le 2 \cdot \left(\frac{e\varepsilon}{50}\right)^{50 \ln n / c_1 \varepsilon} \le n^{-11}$$

for n large enough.

$$\mathbb{E}|\{i : \exists k \in \{0, \dots, k_0\}, \text{ such that } e_{ki}^{row} > c_\varepsilon \cdot 2^{-k} n\}| \le n^{-10},$$

then, for $K = \{0, \dots, k_0\}$,

$$\mathbb{P}\{|R_{heavy}^K| \ge 1\} \le en^{-10}$$

Clearly the same estimate is true for the heavy columns. Finally, with probability at least

$$1 - 2en^{-10} - 2\exp(-n\varepsilon/2) \ge 1 - 6n^{-10}$$

all four good events hold.

Then, note that $\mathbb{P}\{|A_{ij}| > q_{k_1}\} = 2^{-k_1} \le \frac{\varepsilon}{25n}$. Then exactly the same argument as in Lemma 4.21 (the part of proof of Theorem 4.1 for very large entries) implies that with probability at least $1-e^{-\varepsilon n}$ matrix A_3 contains at most εn non-zero entries.

□

The following lemma checks that the matrix A restricted to the small enough entries has a norm of order $\sqrt{n}$:

Lemma 5.16. *Suppose A is a random $n \times n$ matrix with symmetric i.i.d. entries A_{ij} such that $\mathbb{E}A_{ij}^2 = 1$. Let q_k be 2^{-k}-quantiles of the distribution of $|A_{ij}|$ defined by* (5.16)*. Let S contains the entries of A that satisfy $|A_{ij}| \le q_{k_0}$, where $k_0 := \lfloor \log_2 \frac{c_1 n}{\ln n} \rfloor$ with some positive constant c_1 small enough. Then with probability at least $1 - n^{-9}$ we have*

$$\|S\| \le 8\sqrt{n}.$$

Proof. Note that

$$\mathbb{P}\{|A_{ij}| > q_{k_0}\} = 2^{-k_0} \ge \frac{\ln n}{c_1 n} \ge \mathbb{P}\left\{|A_{ij}| > \sqrt{\frac{c_1 n}{\ln n}}\right\},$$

so, all the elements in S are such that $|A_{ij}| \le q_{k_0} \le \sqrt{c_1 n/\ln n}$. Then the norm of S can be estimated by Bandeira-van Handel Theorem 5.2, exactly in the same way it was done in the proof of Theorem 5.1 (with $K := \sqrt{\frac{c_1 n}{\ln n}}$).

The application of Theorem 5.2 with $t = \sqrt{n}$ gives

$$\mathbb{P}\{\|A_1\| \ge 8\sqrt{n}\} \le n \exp(-\frac{c}{4c_1} \ln n) \le n \cdot n^{-10} = n^{-9},$$

if we take any $c_1 \le c/40$ (c here is a constant defined by Theorem 5.2). □

To bound the norm of the rest, we will use several known results for Bernoulli random matrices. The following two lemmas are the versions of [Lemma 22, [CRV15]], in modification from the proof of [Lemma 12, [CRV15]], and [Lemma 22, [CRV15]]. Originally they were proved in [FO05] for symmetric Bernoulli matrices – adjacency matrices of the random graphs – but their proofs follow without modifications for the non-symmetric case.

Lemma 5.17. *Let B be a $n \times n$ Bernoulli 0-1 matrix with $\mathbb{P}\{B_{ij} = 1\} = p$. Let $\tilde{B}$ be the matrix obtained from B by zeroing out some rows and columns of B, such that $\tilde{B}$ has at most $C_0 np$ positive entries in every row and column. Let $e(S,T) := \sum_{i \in S, j \in T} \tilde{B}_{ij}$ (i.e. number of non-zero elements in the submatrix spanned by $S \times T$).*

Then with probability at least $1 - n^{-10}$ for any $S, T \subset [n]$ one the following holds:

(A) $e(S,T) \le C_1 |S||T|p$, *or*

(B) $e(S,T)\cdot\log\left(\frac{e(S,T)}{|S||T|p}\right)\le C_2|T|\log\left(\frac{n}{|T|}\right)$,

where C_1 and C_2 are constants independent from S,T and n. Moreover, C_2 is independent from C_0 and $C_1 = C_0\cdot e$.

The second lemma is a deterministic statement. It shows that bounded rows and columns, and one of the two tail conditions (A) or (B) are enough to bound the quadratic form $\langle Bu, v\rangle$ for all unit vectors u and v:

Lemma 5.18. *Let B be an $n\times n$ matrix with 0-1 elements and $p>0$, such that every its row and column contains at most C_3np ones and for any $S,T\subset[n]$ either condition (A) or (B) (from Lemma 5.17) holds . Then for any $u,v\in S^{n-1}$*

$$\sum_{i,j:|u_iv_j|\ge\sqrt{p/n}} B_{ij}|u_iv_j|\le C_{5.18}\sqrt{np},$$

where $C_{5.18}=C(C_0,C_1,C_2,C_3)$. In particular, $C_{5.18}\le 8e\cdot C_0$.

The next lemma is a version of [Claim 2.7, [FO05]]. Original claim bounds a contribution of the "light couples" to the quadratic form $\sum_{ij\in L}M_{ij}u_iv_j$, where M_{ij} were Bernoulli random variables, $u=(u_i)_{i=1}^n$ and $v=(v_j)_{j=1}^n$ were some fixed unit vectors, and $L:=\{i,j\in[n]:|u_iv_j|\le 1/\sqrt{n}\}$. In our version M_{ij} are no longer Bernoulli, so the definition of the "light couples" would depend on the absolute values $|M_{ij}|$.

The possibility of additional conditioning on an index subset Q will be explored later when we apply Lemma 5.19 to the regularized matrix (with some rows and columns deleted).

Lemma 5.19. *Consider an $n\times n$ random matrix M with independent symmetric entries and $\mathbb{E}M_{ij}^2\le 1$. Consider two vectors $u=(u_i)_{i=1}^n$ and $v=(v_j)_{j=1}^n$ such that $u,v\in S^{n-1}$. Denote the event $\mathcal{M}_{ij}^{light}=\{|M_{ij}||u_iv_j|\le\sqrt{2/n}\}$ and $Q\subset[n]\times[n]$ is an index subset. Then for any constant $C\ge 2$*

$$|\sum_{ij}u_iM_{ij}\mathbb{1}_{\{(i,j)\in Q\}}\mathbb{1}_{\mathcal{M}_{ij}^{light}}v_j|\le C\sqrt{n}$$

with probability at least $1-2\exp(-Cn/2)$.

Proof. Let $R_{ij}:=M_{ij}\mathbb{1}_{\{(i,j)\in Q\}}\mathbb{1}_{\mathcal{M}_{ij}^{light}}$. Note that R_{ij} are centered due to the symmetric distribution of M_{ij}, and they are independent as M_{ij} are. So we can apply Bernstein's inequality for bounded distributions (Lemma 2.7) to bound the sum:

$$\mathbb{P}\{|\sum_{ij} u_i R_{ij} v_j| \geq t\} \leq 2\exp\left(-\frac{t^2/2}{\sigma^2 + Kt/3}\right),$$

where

$$K = \max_{i,j} |u_i R_{ij} v_j| \leq \sqrt{2/n} \quad \text{and} \quad \sigma^2 = \sum_{ij} \mathbb{E}(u_i R_{ij} v_j)^2.$$

Note that $\mathbb{E}R_{ij}^2 \leq \mathbb{E}M_{ij}^2$, as $R_{ij}^2 \leq M_{ij}^2$ almost surely, and $\mathbb{E}M_{ij}^2 \leq 1$. So,

$$\sigma^2 = \sum_{ij} u_i^2 \mathbb{E}R_{ij}^2 v_j^2 \leq \sum_{i,j} u_i^2 v_j^2 = 1,$$

as $\sum_i u_i^2 = \sum_j v_j^2 = 1$. So, taking $t = C\sqrt{n}$, we obtain

$$\mathbb{P}\{|\sum_{(i,j)} u_i M_{ij} \mathbb{1}_{\{(i,j)\in Q\}} \mathbb{1}_{\mathcal{M}_{ij}^{light}} v_j| \geq C\sqrt{n}\} \leq 2\exp(-Cn/2)$$

for any $C \geq 2$. This concludes the statement of the lemma. □

Finally, we will use the following simple lemma about the rate of growth about of the 2^{-k}-quantiles of $|A_{ij}|$.

Lemma 5.20. *For every "level" $k = 0, \dots, k_1$ we have*

$$q_k \leq 2^{(k+1)/2} \tag{5.17}$$

Proof. Indeed, we can trivially estimate

$$A_{ij}^2 \geq \sum_{k=1}^{k_1+1} q_{k-1}^2 \mathbb{1}_{\{|A_{ij}|\in(q_{k-1},q_k]\}}.$$

Hence, due to the moment condition $\mathbb{E}A_{ij}^2 = 1$,

$$1 = \mathbb{E}A_{ij}^2 \geq \sum_{k=1}^{k_1+1} q_{k-1}^2 2^{-k} = \sum_{k=0}^{k_1} q_k^2 2^{-k-1}.$$

Multiplying each side by to, we have an estimate for the sum

$$\sum_{k=0}^{k_1} q_k^2 2^{-k} \leq 2. \tag{5.18}$$

As all the term are non-negative, for any $k = 0, \dots, k_1$ we have

$$q_k 2^{-k/2} \le \sqrt{\sum_{k=0}^{k_1} q_k^2 2^{-k}} \le \sqrt{2}.$$

This proves Lemma 5.20. $\square$

Now we are ready to combine these results to prove the key part of Theorem 5.14: the bound for the norm $\tilde{M} = \tilde{A} \cdot \mathbb{1}_{\{q_{k_0} < |A_{ij}| \le q_{k_1}\}}$.

Proposition 5.21. *Suppose A is a random $n \times n$ matrix with i.i.d. symmetric entries A_{ij} with $\mathbb{E}A_{ij}^2 = 1$. Let q_k be 2^{-k}-quantiles defined by (5.16) and let $\tilde{A}$ be its regularized version after we deleted rows and columns that contain more than $(c_\varepsilon 2^{-k} n)$ elements such that $q_{k-1} \le |A_{ij}| < q_k$ for some $k \in \{k_0 + 1, \dots, k_1\}$. Here $c_\varepsilon = 50/\varepsilon$, $k_0 := \lfloor \log_2 \frac{c_1 n}{\ln n} \rfloor$ (for c_1 defined in Lemma 5.16) and $k_1 = \lceil \log_2 \frac{25n}{\varepsilon} \rceil$. Let*

$$\tilde{M} := \tilde{A} \cdot \mathbb{1}_{\{|\tilde{A}_{ij}| \in (q_{k_0}, q_{k_1}]\}}.$$

Then with probability at least $1 - 2n^{-9}$ we have

$$\|\tilde{M}\| \le C_\varepsilon \sqrt{n \ln \ln n}.$$

Here, $C\varepsilon = C \cdot c_\varepsilon$ where C is an absolute constant.

Proof. **Step 1. Net approximation.**

Let $\mathcal{N}$ be a $1/2$-net on S^{n-1} with cardinality $|\mathcal{N}| \le 5^n$ (the existence of such net is a standard fact that can be found, e.g. in [Ver16]). We will use a simple net approximation of the norm (see, e.g. [Lemma 4.4.1, [Ver16]]), namely,

$$\|\tilde{M}\| \le 4 \max_{u,v \in \mathcal{N}} \langle \tilde{M} u, v \rangle = 4 \max_{u,v \in \mathcal{N}} |\sum_{ij} \tilde{M}_{ij} u_i v_j|.$$

We will split the sum into two parts and bound each of them separately (this part is a variant of Feige and Ofek argument, presented in [FO05, CRV15]), based on the absolute value of the element. Let $M := A \cdot \mathbb{1}_{\{|A_{ij}| \in (q_{k_0}, q_{k_1}]\}}$.

For any fixed $u, v \in \mathcal{N}$ and every $i, j \in [n]$ we can define an event

$$\mathcal{M}_{ij}^{light} = \{|M_{ij}||u_i v_j| \ge \sqrt{2/n}\}.$$

Then,

$$
\begin{aligned}
&\max_{u,v\in\mathcal{N}}|\sum_{ij}\tilde{M}_{ij}u_iv_j| \\
&\qquad \le \max_{u,v\in\mathcal{N}}|\sum_{ij}\tilde{M}_{ij}(\mathbb{1}_{\mathcal{M}_{ij}^{light}}+\mathbb{1}_{(\mathcal{M}_{ij}^{light})^c})u_iv_j| \\
&\qquad \le \max_{u,v\in\mathcal{N}}|\sum_{ij}\tilde{M}_{ij}\mathbb{1}_{\mathcal{M}_{ij}^{light}}u_iv_j| + \max_{u,v\in\mathcal{N}}|\sum_{ij}\tilde{M}_{ij}\mathbb{1}_{(\mathcal{M}_{ij}^{light})^c}u_iv_j|.
\end{aligned}
$$

Step 2. Light members.

By Lemma 5.19, for any fixed $u, v \in S^{n-1}$ and a fixed subset of indices Q (assuming that Q^c is a set of rows and columns to delete),

$$|\sum_{ij}u_iM_{ij}\mathbb{1}_{\{(i,j)\in Q\}}\mathbb{1}_{\mathcal{M}_{ij}^{light}}v_j| > 12\sqrt{n} \tag{5.19}$$

with probability at most $2\exp(-6n)$. Now, taking union bound over 5^n choices for u, as many choices for v, and 2^{2n} choices for the row and column subset Q^c, we obtain that

$$\mathbb{P}\{|\sum_{ij}u_i\tilde{M}_{ij}\mathbb{1}_{\mathcal{M}_{ij}^{light}}v_j| \le 12\sqrt{n}\} \ge 1-2\exp(-n). \tag{5.20}$$

Step 3. Other members.

The second term can be roughly bounded by the sum of absolute values:

$$
\begin{aligned}
&|\sum_{ij}\tilde{M}_{ij}\mathbb{1}_{(\mathcal{M}_{ij}^{light})^c}u_iv_j| \\
&\qquad \le \sum_{ij}|\tilde{M}_{ij}|\mathbb{1}_{(\mathcal{M}_{ij}^{light})^c}|u_iv_j| \\
&\qquad \le \sum_{ij}(\sum_{k=k_0+1}^{k_1}q_k\mathbb{1}_{\{|\tilde{M}_{ij}|\in(q_{k-1},q_k]\}})\mathbb{1}_{\{|M_{ij}||u_iv_j|\ge\sqrt{2/n}\}}|u_iv_j|
\end{aligned}
$$

Note that as long as $\mathbb{1}_{\{|\tilde{M}_{ij}|\in(q_{k-1},q_k]\}} = 1$ we also have that $|M_{ij}| \le q_k$. Indeed, $|M_{ij}| > q_k$ implies either $|\tilde{M}_{ij}| > q_k$ or $|\tilde{M}_{ij}| = 0$. In any case, $|\tilde{M}_{ij}| \notin (q_{k-1}, q_k]$. So, the last expression is bounded above by

$$\sum_{ij}\sum_{k=k_0+1}^{k_1}q_k\mathbb{1}_{\{|\tilde{M}_{ij}|\in(q_{k-1},q_k]\}}\mathbb{1}_{\{q_k|u_iv_j|\ge\sqrt{2/n}\}}|u_iv_j|$$

Using Lemma 5.20, we can further estimate

$$\mathbb{1}_{\{q_k|u_iv_j|\geq\sqrt{2/n}\}} \leq \mathbb{1}_{\{2^{(k+1)/2}|u_iv_j|\geq\sqrt{2/n}\}} = \mathbb{1}_{\{|u_iv_j|\geq\sqrt{2^{-k}/n}\}}.$$

As a result, we got

$$|\sum_{ij} \tilde{M}_{ij}\mathbb{1}_{(\mathcal{M}_{ij}^{light})^c}u_iv_j| \leq \sum_{k=k_0+1}^{k_1} q_k \sum_{ij:|u_iv_j|\geq\sqrt{2^{-k/n}}} \mathbb{1}_{\{|\tilde{M}_{ij}|\in(q_{k-1},q_k]\}}|u_iv_j|. \tag{5.21}$$

Step 4. Bernoulli matrices. For each "level" $k = k_0+1,\dots,k_1$ let us define a matrix

$$B^k = (B^k_{ij})^n_{i,j=1} := \mathbb{1}_{\{|M_{ij}|\in(q_{k-1},q_k]\}}$$

It has independent Bernoulli entries with $\mathbb{E}B^k_{ij} = 2^{-k}$. Note that the matrix

$$\tilde{B}^k = (\tilde{B}^k_{ij})^n_{i,j=1} := \mathbb{1}_{\{|\tilde{M}_{ij}|\in(q_{k-1},q_k]\}}$$

is obtained from B by zeroing out some rows and columns, such that new maximal degree is at most $c_\varepsilon n 2^{-k}$. So, we can apply Lemma 5.17 to $\tilde{B}_k$ to conclude that one of the conditions on the number of ones in sub-blocks (A) or (B) holds with probability at least n^{-10}.

So, with probability at least $1-n^{-10}|k_1-k_0|$ we can apply Lemma 5.18 to every $\tilde{B}_k$ with $k = k_0+1,\dots,k_1$. Which means that by Lemma 5.18,

$$\sum_{ij:|u_iv_j|\geq\sqrt{2^{-k/n}}} \tilde{B}^k_{ij}|u_iv_j| \leq C_{5.18}\sqrt{2^{-k}n} \quad \text{for every } k = k_0+1,\dots,k_1. \tag{5.22}$$

Step 5. Conclusion. Note that the number of terms in the sum

$$|k_1-k_0| \leq \log_2\frac{25n}{\varepsilon} - \log_2\frac{c_1 n}{\ln n} + 2 \leq 2\log_2\ln n \tag{5.23}$$

for all large enough n.

So, we can combine the estimates (5.21) and (5.22) to claim that with probability at least $1-2n^{-10}\log_2\ln n \geq 1-n^{-9}$

$$|\sum_{ij} \tilde{M}_{ij}\mathbb{1}_{(\mathcal{M}_{ij}^{light})^c}u_iv_j| \leq C_{5.18}\sum_{k=k_0}^{k_1} q_k\sqrt{2^{-k}n}.$$

To estimate right hand side, recall that $C_{5.18} \leq 8ec_\varepsilon$. Also, by Cauchy-Schwarz, (5.18) and the number of terms estimate (5.23), we have that

$$\sum_{k=k_0}^{k_1} q_k\sqrt{2^{-k}} \le \sqrt{\sum_{k=k_0}^{k_1} q_k^2 2^{-k}}\sqrt{\sum_{k=k_0}^{k_1} 1} \le 4\sqrt{\ln\ln n}.$$

Finally, combining this with the estimate for the light part (5.20), we get

$$\|\tilde{M}\| \le 12\sqrt{n} + 4C_{5.18}\sqrt{n\cdot\ln\ln n} \lesssim c_\varepsilon\sqrt{n\cdot\ln\ln n}$$

with probability at least $1-2e^{-n}-n^{-9} \ge 1-2n^{-9}$ for all n large enough.

$\square$

Proof of Theorem 5.14. Decompose A into a sum of three $n\times n$ matrices with disjoint support,

$$A = S + M + L,$$

where S contains the entries of A that satisfy $|A_{ij}| \le q_{k_0}$, the matrix M contains the entries for which $q_{k_0} < |A_{ij}| \le q_{k_1}$, and L contains large entries – those satisfying $|A_{ij}| > q_{k_1}$. Here, q_{k_0} and q_{k_1} are quantiles of the distribution of A_{ij}, defined by (5.16), and $k_0 := \lfloor \log_2 \frac{c_1 n}{\ln n}\rfloor$ (with a small constant c_1 defined in the Lemma 5.16), $k_1 = \lceil \log_2 \frac{25n}{\varepsilon}\rceil$.

By Lemma 5.15, there is an event $\mathcal{E}$ (with $\mathbb{P}\{\mathcal{E}\} \ge 1-6n^{-10}$), on which a) the regularization procedure of Theorem 5.14 is local (changes at most εn rows, εn columns and εn entries of A) and b) S does not have rows or columns that are too "heavy".

Hence, when $\mathcal{E}$ holds, the regularization does not touch the elements of S, deletes the heavy rows and columns in M (let us call the resulting matrix $\tilde{M}$) and completely zeroes out L, as all its non-zero entries are greater than q_{k_1}. So,

$$\tilde{A} = S + \tilde{M},$$

Now we are going to estimate the norms of S and $\tilde{M}$ separately. Elements of S are independent (since they are obtained by independent individual truncations from the i.i.d. elements A_{ij}) and have zero mean (since A_{ij} have symmetric distribution). So, we can apply Lemma 5.16 to conclude that

$$\mathbb{P}\{\|S\| > 8\sqrt{n}\} \le n^{-9}.$$

Finally, the norm of $\tilde{M}$ is estimated in Proposition 5.21. We proved that

$$\mathbb{P}\{\|\tilde{A}_2\| > Cc_\varepsilon\sqrt{n\cdot\ln\ln n}\} \le 2n^{-9}.$$

Hence, we have that

$$\|\tilde{A}\| \lesssim c_\varepsilon \sqrt{n \cdot \ln \ln n}$$

and the regularization process described in Theorem 5.14 is local with probability at least $1 - 6n^{-10} - n^{-9} - 2n^{-9} \geq 1 - 9n^{-9}$. Theorem 5.14 is proved. $\square$

5.3.2 Discussion of Theorem 5.14

Dependence on ε

Theorem 5.14 shows that

$$\|\tilde{A}\| \leq C c_\varepsilon \sqrt{n \cdot \ln \ln n}.$$

with $c_\varepsilon = 50/\varepsilon$. Thus, the dependence of the resulting matrix norm on the fraction of the entries to be deleted is of order $1/\varepsilon$. This is worse than order $\ln \varepsilon^{-1}/\sqrt{\varepsilon}$ dependence (that we obtained in Theorem 4.1), and, of course, $1/\sqrt{\varepsilon}$ (that we conjectured in Remark 4.2 as optimal).

However, with a slightly more complex version the algorithm presented in Theorem 5.14 the norm can be locally regularized to the order $c'_\varepsilon \sqrt{\log \log n \cdot n}$ with

$$c'_\varepsilon = \frac{\ln(\varepsilon^{-1})}{\sqrt{\varepsilon}}.$$

Namely, we could do the same rows and columns regularization as described in Theorem 5.14 with $c_\varepsilon \sim \ln(\varepsilon^{-1})$, only up to the level q_{k_2} with $k_2 = \lceil \log_2 \frac{n}{4} \rceil$. For the entries A_{ij} such that $|A_{ij}| \in (q_{k_2}, q_{k_1}]$ we would need to do a separate regularization procedure, like the one we used in Section 4.2.3. Considering the rows of A with more than $c2^{-k_2}n + 2\ln \varepsilon^{-1}$ elements such that $q_{k_2} < |A_{ij}| \leq q_{k_1}$, we can prove that with highh probability they do not have more than $O(\varepsilon n)$ such elements collectively. As in Lemma 4.17, deletion of all such elements achieves the norm of order $\sqrt{n} \ln(\varepsilon^{-1})/\sqrt{\varepsilon}$.

We will still have extra $\sqrt{\ln \ln n}$ factor from the estimate for A_{ij} such that $|A_{ij}| \leq q_{k_2}$. So, we decided to concentrate on simpler regularization algorithm in Theorem 5.14, as the ε-dependence loss is still negligible comparing with the extra $\sqrt{\ln \ln n}$ term.

Dependence on n

The main disadvantage of our result is, of course, that the resulting norm $\|\tilde{A}\| \gg \sqrt{n}$. The reason for the extra n-dependent term is that in the proof we consider restrictions of A to the discretization "levels" independently, and independently estimate their norms. The second moment assumption gives us that $\sum q_k^2 2^{-k} \sim 1$. However, the best we can hope for a norm of one "level" (after proper regularization) is $q_k\sqrt{n2^{-k}}$ (since this is an expected L_2-norm of a restricted row). Thus, we end up summing square roots of the converging series, $\sum q_k 2^{-k/2}$, which can be as large as square root of the number of summands ($\ln\ln n$ in our case).

Symmetry assumption

Another additional condition we have is symmetric distribution of the entries. We use this to keep zero mean after various truncations by absolute value. A simple argument that helped us before to claim that the mean cannot change too much in truncation (e.g. in the proof of Theorems 4.1 and 5.1) does not work when the level of truncation is low enough (like we have in Lemmas 5.16 or 5.19). The standard symmetrization techniques also would not work since we combine the convex norm function with truncation (zeroing out of some rows), which is not convex. But it is likely that some regularization procedure (in flavor of the one we used in Theorem 3.11 and 4.13 for $\|.\|_{\infty\to 2}$) would work for the Theorem 5.14.

Conjectured optimal regularization

Finally, what would be an ideal conjectured way to locally regularize the norm? Clearly, we have to delete all the rows and columns with L_2-norm greater than $\sqrt{n}$. Is this enough?

A result by Seginer ([Seg00]) shows that in expectation the norm of the matrix with i.i.d. elements is bounded by the largest norm of its row or column. However, note that after cutting "heavy" rows and columns we lose independence of the entries in the resulting matrix, so the result of Seginer cannot be directly applied.

It would be desirable to remove extra $\sqrt{\log\log n}$ term, and to simplify the regularization procedure in the Theorem 5.14, proving something like the following

Conjecture 5.22. *Consider an $n\times n$ random matrix A with i.i.d. mean zero variance one elements. Let $R_i(A)$ and $T_i(A)$ denote the i-th row and column vector of the matrix A respectively. Let $\tilde{A}$ be the matrix that obtained from A by zeroing out all rows and columns such that*

$$\|R_i(A)\|_m \geq C\mathbb{E}\|R_i(A)\|_m, \quad \|T_i(A)\|_m \geq C\mathbb{E}\|T_i(A)\|_m \tag{5.24}$$

for some L_m-norm to be specified (e.g. $m=2$).

Then with probability $1-o(1)$ the operator norm satisfies $\|\tilde{A}\| \leq C'\sqrt{n}$.

CHAPTER 6

Further directions

In the end of the Chapter 5 we discussed an open question regarding constructive regularization of the norms of the square matrices with i.i.d. entries having exactly two finite moments. In this chapter we would like to collect several more open questions related to the work described in this thesis.

• **Optimal moment assumptions.** The key step of our proof of the heavy-tailed invertibility theorem (Theorem 3.2 – construction of the universal covering) requires finite second moment assumption. But it is an open question to identify an optimal moment condition for the invertibility estimate (1.5) for square random matrices having i.i.d. elements.

For example, in the case of random rectangular matrices prior work [LPRTJ05, RV09] was extended by K. Tikhomirov ([Tik16]) to the matrices with no assumptions on moments at all. However, this argument works only for the matrices of the size $m \times n$ such that $m \leq \delta n$ for some constant $\delta < 1$.

• **Intermediate singular values.** The work discussed above, in particular, answers the question if s_{min} and s_{max} typically have the same order as in "ideal" gaussian case: for $s_{min} = s_n$ the answer is "yes," and for $s_{max} = s_1$ the answer is "no" (but local regularization is possible).

It would be interesting to know the answer for the other s_k, $1 < k < n$. The distribution of all s_k for "ideal" square matrices is known: $s_k \sim \frac{n+1-k}{\sqrt{n}}$ (see [Sza90] for the Gaussian case, [RV09] and [Wei17] imply the same estimate for subgaussian under an additional mild assumption). Is there a threshold k_0, such that for $k > k_0$ the order of s_k stays the same for the heavy-tailed matrices, and for $k < k_0$ it is not? How does the size of the submatrix to be deleted depend on $k > 1$? One way to approach this question is try to combine refined net and regularization techniques developed in [RV18, RT15] with the analysis of the projections onto subspaces in the flavor of [Wei17].

• **Non-identically distributed entries.** We use the i.i.d. entries assumption in Theorem 1.1 several times, in particular, adopting from [RV08] the small ball probability estimate via LCD (least common denominator is a notion that measures the additive structure of the vector). One approach to relax identical distribution condition is is to generalize the notion of LCD of a vector to the several not identically distributed vectors, and prove a new small ball probability estimate with it. For example, this can be done in a simple partial case of several Bernoulli vectors.

The lower bound for s_{min} for the matrices with the non-identically distributed entries would help, in particular, to improve probability estimates for the sparse random matrices (in continuation of the work [LR11]).

• **Restricted singular values** Another objects to look at are restricted singular values. For example, smallest singular value of an $m \times n$ matrix A restricted on a closed $T \subset \mathbb{R}^n$ is defined geometrically as

$$s_{min}(A;T) := \inf_{x \in T} \|Ax\|_2.$$

In their recent paper Tropp and Oymak ([OT15]) employ Lindeberg's technique ([Lin22]) to bound a restricted smallest singular value of the matrices with no identical distribution assumed:

$$|s_{min}(A;T) - (\mathcal{E}_m(T))_+| \ll \sqrt{m+n}$$

with high probability. Here, $\mathcal{E}_m(T)$ is a m-excess width of the set T:

$$\mathcal{E}_m(T) := \mathbb{E} \min_{t \in T}(\sqrt{m} + g \cdot t/\|t\|),$$

where $g \sim N(0, I)$.

This work requires at least fourth finite moment of the entries A_{ij} as well as independence of the entries. It is interesting to relax moment conditions required to make a high probability estimate, as well as allow some dependence (for example, to assume isotropic rows).

• **Markov matrices** An interesting matrix associated with the Erdös-Rényi random graph is Markov kernel matrix that was recently studied in the works of Chafaï, Bordenave, Caputo, and Piras (see [BCCP17, BCC12]).

Markov matrix is a transition probability matrix for a random walk on the underlying graph, for example, an Erdös-Rényi random graph. A natural way to define it is to take $(X_{jk})_{j,k>1}$ independent identically distributed non-negative random variables (under certain regularity conditions), then M is the $n \times n$ random Markov matrix with i.i.d. rows defined by $M_{jk} = X_{jk}/(X_{j1} + \ldots + X_{jn})$. The paper [BCC12] studies the properties of the limiting empirical measure of the eigenvalues of M, which is, informally, characterizes the behavior of the random walk in the long run. The results obtained are based on a non-asymptotic estimate for N:

$$\mathbb{P}\{s_{\min}(\sqrt{n}M - z\,\mathrm{Id}) \le n^{-b}\} \le n^{-a}, \tag{6.1}$$

where $|z| \le C$ and $a, b, C > 0$ are the absolute constants. This is an invertibility estimate of the type obtained in Theorem 3.1, valid for a broader class of matrices (with dependencies and shift). At the moment, (6.1) is proved under condition that elements M_{ij} have bounded density, which automatically implies that underlying Erdös-Rényi graph has to be complete. The authors believe that a similar estimate should hold without boundedness condition, and clearly it would be good to extend the results for the case of non-complete (potentially, sparse) random graphs.

BIBLIOGRAPHY

[AGZ10] Greg W Anderson, Alice Guionnet, and Ofer Zeitouni, *An introduction to random matrices*, vol. 118, Cambridge University Press, Cambridge New York, 2010.

[ALPTJ10] Radosław Adamczak, Alexander Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann, *Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles*, Journal of the American Mathematical Society **23** (2010), no. 2, 535–561.

[ALPTJ11] Radosław Adamczak, Alexander E Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann, *Sharp bounds on the rate of convergence of the empirical covariance matrix*, Comptes Rendus Mathematique **349** (2011), no. 3-4, 195–200.

[AN04] Noga Alon and Assaf Naor, *Approximating the cut-norm via Grothendieck's inequality*, Proceedings of the thirty-sixth annual ACM symposium on Theory of computing, ACM, 2004, pp. 72–80.

[AT16] Antonio Auffinger and Si Tang, *Extreme eigenvalues of sparse, heavy tailed random matrices*, Stochastic Processes and their Applications **126** (2016), no. 11, 3310–3330.

[BCC12] Charles Bordenave, Pietro Caputo, and Djalil Chafaï, *Circular law theorem for random markov matrices*, Probability Theory and Related Fields **152** (2012), no. 3-4, 751–779.

[BCCP17] Charles Bordenave, Pietro Caputo, Djalil Chafaï, and Daniele Piras, *Spectrum of large random markov chains: heavy-tailed weights on the oriented complete graph*, Random Matrices: Theory and Applications **6** (2017), no. 02, 1750006.

[BJR10] B Bollobás, S Janson, and O Riordan, *The cut metric, random graphs, and branching processes, preprint (2009)*, Journal of statistical physics **140** (2010), no. 2, 289–335.

[Bor17] Gaëtan Borot, *An introduction to random matrix theory*, arXiv preprint arXiv:1710.10792 (2017).

[BSY88] Zhidong D Bai, Jack W Silverstein, and Yong Q Yin, *A note on the largest eigenvalue of a large dimensional sample covariance matrix*, Journal of Multivariate Analysis **26** (1988), no. 2, 166–168.

[BVH16] Afonso S Bandeira and Ramon Van Handel, *Sharp nonasymptotic bounds on the norm of random matrices with independent entries*, The Annals of Probability **44** (2016), no. 4, 2479–2506.

[BY08] ZD Bai and YQ Yin, *Limit of the smallest eigenvalue of a large dimensional sample covariance matrix*, Advances In Statistics, World Scientific, 2008, pp. 108–127.

[CRV15] Peter Chin, Anup Rao, and Van Vu, *Stochastic block model and community detection in sparse graphs: A spectral algorithm with optimal rate of recovery*, Conference on Learning Theory, 2015, pp. 391–423.

[Ede88] Alan Edelman, *Eigenvalues and condition numbers of random matrices*, SIAM Journal on Matrix Analysis and Applications **9** (1988), no. 4, 543–560.

[ESY09] László Erdős, Benjamin Schlein, and Horng-Tzer Yau, *Local semicircle law and complete delocalization for Wigner random matrices*, Communications in Mathematical Physics **287** (2009), no. 2, 641–655.

[ESY10] ______, *Wegner estimate and level repulsion for Wigner random matrices*, International Mathematics Research Notices **2010** (2010), no. 3, 436–479.

[FO05] Uriel Feige and Eran Ofek, *Spectral techniques applied to sparse random graphs*, Random Structures & Algorithms **27** (2005), no. 2, 251–275.

[GLPTJ17] Olivier Guédon, Alexander E Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann, *On the interval of fluctuation of the singular values of random matrices*, Journal of the European Mathematical Society **19** (2017), 1469–1505.

[GNT15] Friedrich Götze, Alexey Naumov, and Alexander Tikhomirov, *On minimal singular values of random matrices with correlated entries*, Random Matrices: Theory and Applications **4** (2015), no. 02, 1550006.

[Hoe63] Wassily Hoeffding, *Probability inequalities for sums of bounded random variables*, Journal of the American statistical association **58** (1963), no. 301, 13–30.

[KM15] Vladimir Koltchinskii and Shahar Mendelson, *Bounding the smallest singular value of a random matrix without concentration*, International Mathematics Research Notices **2015** (2015), no. 23, 12991–13008.

[Koz12] Dexter C Kozen, *The design and analysis of algorithms*, Springer Science & Business Media, 2012.

[Lat05] Rafał Latała, *Some estimates of norms of random matrices*, Proceedings of the American Mathematical Society **133** (2005), no. 5, 1273–1282.

[Lin22] Jarl Waldemar Lindeberg, *Eine neue herleitung des exponentialgesetzes in der wahrscheinlichkeitsrechnung*, Mathematische Zeitschrift **15** (1922), no. 1, 211–225.

[LLV17] Can M Le, Elizaveta Levina, and Roman Vershynin, *Concentration and regularization of random graphs*, Random Structures & Algorithms **51** (2017), no. 3, 538–561.

[LLV18a] ———, *Concentration of random graphs and application to community detection*, arXiv preprint arXiv:1801.08724 (2018).

[LLV18b] ———, *Concentration of random graphs and application to community detection*, arXiv preprint arXiv:1801.08724 (2018).

[LPRTJ05] Alexander E Litvak, Alain Pajor, Mark Rudelson, and Nicole Tomczak-Jaegermann, *Smallest singular value of random matrices and geometry of random polytopes*, Advances in Mathematics **195** (2005), no. 2, 491–523.

[LR11] Alexander Litvak and Omar Rivasplata, *Smallest singular value of sparse random matrices*, arXiv preprint arXiv:1106.0938 (2011).

[LR15] Jing Lei and Alessandr Rinaldo, *Consistency of spectral clustering in stochastic block models*, The Annals of Statistics **43** (2015), no. 1, 215–237.

[LS14] Alexander E Litvak and Susanna Spektor, *Quantitative version of a silverstein's result*, Geometric Aspects of Functional Analysis, Springer, 2014, pp. 335–340.

[LT13] Michel Ledoux and Michel Talagrand, *Probability in Banach spaces: isoperimetry and processes*, Springer Science & Business Media, 2013.

[MP14] Shahar Mendelson and Grigoris Paouris, *On the singular values of random matrices*, Preprint (2014).

[MS86] Vitali D Milman and Gideon Schechtman, *Asymptotic theory of finite dimensional normed spaces: Isoperimetric inequalities in riemannian manifolds*, vol. 1200, Springer Science & Business Media, 1986.

[NO14] Hoi H Nguyen and Sean O'Rourke, *The elliptic law*, International Mathematics Research Notices **2015** (2014), no. 17, 7620–7689.

[OT15] Samet Oymak and Joel A Tropp, *Universality laws for randomized dimension reduction, with applications*, Information and Inference: A Journal of the IMA (2015).

[Pie78] Albrecht Pietsch, *Operator ideals*, vol. 16, Deutscher Verlag der Wissenschaften, 1978.

[Pis99] Gilles Pisier, *The volume of convex bodies and Banach space geometry*, vol. 94, Cambridge University Press, 1999.

[Rog61] Boris Alekseevich Rogozin, *On the increase of dispersion of sums of independent random variables*, Theory of Probability & Its Applications **6** (1961), no. 1, 97–99.

[RT15] Elizaveta Rebrova and Konstantin Tikhomirov, *Coverings of random ellipsoids, and invertibility of matrices with iid heavy-tailed entries*, arXiv preprint arXiv:1508.06690 (2015).

[Rud08] Mark Rudelson, *Invertibility of random matrices: norm of the inverse*, Annals of Mathematics (2008), 575–600.

[Rud13] ______, *Lecture notes on non-aymptotic random matrix theory*, arXiv preprint arXiv:1301.2382 **8** (2013).

[RV08] Mark Rudelson and Roman Vershynin, *The LittlewoodOfford problem and invertibility of random matrices*, Advances in Mathematics **218** (2008), no. 2, 600–633.

[RV09] ______, *Smallest singular value of a random rectangular matrix*, Communications on Pure and Applied Mathematics **62** (2009), no. 12, 1707–1739.

[RV10] ______, *Non-asymptotic theory of random matrices: extreme singular values*, Proceedings of the International Congress of Mathematicians 2010 (ICM 2010) (In 4 Volumes) Vol. I: Plenary Lectures and Ceremonies Vols. II–IV: Invited Lectures, World Scientific, 2010, pp. 1576–1602.

[RV18] Elizaveta Rebrova and Roman Vershynin, *Norms of random matrices: local and global problems*, Advances in Mathematics **324** (2018), 40–83.

[Sch11] Jssai Schur, *Bemerkungen zur theorie der beschränkten bilinearformen mit unendlich vielen veränderlichen.*, Journal für die reine und Angewandte Mathematik **140** (1911), 1–28.

[Seg00] Yoav Seginer, *The expected norm of random matrices*, Combinatorics, Probability and Computing **9** (2000), no. 2, 149–166.

[Sil89] Jack W Silverstein, *On the weak limit of the largest eigenvalue of a large dimensional sample covariance matrix*, Journal of Multivariate Analysis **30** (1989), no. 2, 307–311.

[SV13] Nikhil Srivastava and Roman Vershynin, *Covariance estimation for distributions with* $\{2+\varepsilon\}$*-moments*, The Annals of Probability **41** (2013), no. 5, 3081–3111.

[Sza90] Stanislaw J Szarek, *Spaces with large distance to* l_∞^n *and random matrices*, American Journal of Mathematics **112** (1990), no. 6, 899–942.

[Sza91] ______, *Condition numbers of random matrices*, Journal of Complexity **7** (1991), no. 2, 131–149.

[Tao12] Terence Tao, *Topics in random matrix theory*, vol. 132, American Mathematical Soc., 2012.

[Tho00] Hermann Thorisson, *Coupling, stationarity, and regeneration*, Springer, 2000.

[Tik15] Konstantin Tikhomirov, *The limit of the smallest singular value of random matrices with iid entries*, Advances in Mathematics **284** (2015), 1–20.

[Tik16] Konstantin E Tikhomirov, *The smallest singular value of random rectangular matrices with no moment assumptions on entries*, Israel Journal of Mathematics **212** (2016), no. 1, 289–314.

[Tik17] Konstantin Tikhomirov, *Invertibility via distance for non-centered random matrices with continuous distributions*, arXiv preprint arXiv:1707.09656 (2017).

[Tro15] Joel A Tropp, *An introduction to matrix concentration inequalities*, Foundations and Trends in Machine Learning **8** (2015), no. 1-2, 1–230.

[TV08] Terence Tao and Van Vu, *Random matrices: the circular law*, Communications in Contemporary Mathematics **10** (2008), no. 02, 261–307.

[TV09] Terence Tao and Van H Vu, *Inverse Littlewood-Offord theorems and the condition number of random discrete matrices*, Annals of Mathematics (2009), 595–632.

[TV10a] Terence Tao and Van Vu, *Random matrices: The distribution of the smallest singular values*, Geometric And Functional Analysis **20** (2010), no. 1, 260–297.

[TV10b] ———, *Smooth analysis of the condition number and the least singular value*, Mathematics of Computation **79** (2010), no. 272, 2333–2352.

[Ver12] Roman Vershynin, *Introduction to the non-asymptotic analysis of random matrices*, Compressed sensing: Theory and Applications (2012), 210–268.

[Ver16] ———, *High-dimensional probability*, An Introduction with Applications (2016).

[vH17a] Ramon van Handel, *On the spectral norm of gaussian random matrices*, Transactions of the American Mathematical Society **369** (2017), no. 11, 8161–8178.

[vH17b] ———, *Structured random matrices*, Convexity and Concentration, Springer, 2017, pp. 107–156.

[Vu05] Van H Vu, *Spectral norm of random matrices*, Proceedings of the thirty-seventh annual ACM symposium on Theory of computing, ACM, 2005, pp. 423–430.

[Wei17] Feng Wei, *Upper bound for intermediate singular values of random matrices*, Journal of Mathematical Analysis and Applications **445** (2017), no. 2, 1530–1547.

[Wol99] Elmar Wolfstetter, *Topics in microeconomics: Industrial organization, auctions, and incentives*, Cambridge University Press, 1999.

[Yas14] Pavel Yaskov, *Lower bounds on the smallest eigenvalue of a sample covariance matrix.*, Electronic Communications in Probability **19** (2014).

[YBK88] Yong-Qua Yin, Zhi-Dong Bai, and Pathak R Krishnaiah, *On the limit of the largest eigenvalue of the large dimensional sample covariance matrix*, Probability theory and related fields **78** (1988), no. 4, 509–521.