OPTIMAL AND ALGORITHMIC NORM REGULARIZATION OF RANDOM MATRICES

VISHESH JAIN, ASHWIN SAH, AND MEHTAAB SAWHNEY

ABSTRACT. Let A be an $n \times n$ random matrix whose entries are i.i.d. with mean 0 and variance 1. We present a deterministic polynomial time algorithm which, with probability at least $1 - 2\exp(-\Omega(\epsilon n))$ in the choice of A, finds an $\epsilon n \times \epsilon n$ sub-matrix such that zeroing it out results in $\widetilde{A}$ with
$$\|\widetilde{A}\| = O(\sqrt{n/\epsilon}).$$
Our result is optimal up to a constant factor and improves previous results of Rebrova and Vershynin, and Rebrova. We also prove an analogous result for A a symmetric $n \times n$ random matrix whose upper-diagonal entries are i.i.d. with mean 0 and variance 1.

1. INTRODUCTION

Recall that the operator norm of an $n \times n$ real-valued matrix A is defined as
$$\|A\| := \sup_{x \in \mathbb{S}^{n-1}} \|Ax\|_2,$$

where $\|\cdot\|_2$ denotes the Euclidean norm and $\mathbb{S}^{n-1}$ denotes the unit sphere in $\mathbb{R}^n$. The operator norm is a fundamental quantity of interest in the non-asymptotic theory of random matrices (see, e.g., [8] and the references therein). A classical result of Bai, Krishnaiah, and Yin [12] shows that if the entries of an $n \times n$ random matrix A are i.i.d. random variables with 0 mean, unit variance, and bounded fourth moment, then
$$\|A\| = (2 + o(1))\sqrt{n}.$$

The finite fourth moment hypothesis is sharp in the sense that for a sequence of $n \times n$ random matrices A_n with entries that are i.i.d. random variables with mean 0, unit variance, and infinite fourth moment, Bai, Silverstein, and Yin [12] showed that
$$\limsup \frac{\|A\|}{\sqrt{n}} = \infty \text{ almost surely.}$$

Motivated by works of Feige and Ofek [2] and Le, Levina, and Vershynin [4] on the regularization of the norm of adjacency matrices of random graphs, Rebrova and Vershynin [7] asked whether enforcing the bound $\|A\| = O(\sqrt{n})$ is a "local problem" or a "global problem". Specifically, they considered an $n \times n$ random matrix A with i.i.d. entries and asked what assumptions (if any) on the distribution of the entries guarantees that, with high probability, $\|\widetilde{A}\| = O(\sqrt{n})$ for some matrix $\widetilde{A}$ obtained by modifying A on a small sub-matrix. They showed [7, Theorem 1.3] that this is not possible if the distribution has either non-zero mean or infinite variance; in other words, in this case, there is a "global problem". On the other hand, they showed that if the distribution has zero mean and bounded variance, then the problem is "local".

Theorem 1.1 ([7, Theorem 1.1]). *Consider an $n \times n$ random matrix A with i.i.d. entries that have zero mean and unit variance. There exist absolute constants $C_{1.1}, c_{1.1} > 0$ such that for any $\epsilon \in (0, 1/2]$, with probability at least $1 - 2\exp(-c_{1.1}\epsilon n)$, there exists an $\epsilon n \times \epsilon n$ sub-matrix of A*

such that replacing all of its entries with zero gives a matrix $\widetilde{A}$ with

$$\|\widetilde{A}\| \leq C_{1.1} \frac{\log \epsilon^{-1}}{\sqrt{\epsilon}} \sqrt{n}.$$

The work of Rebrova and Vershynin [7] leaves open several natural questions.

1.1. **Optimal norm regularization.** Let $\epsilon \in (0, 1/10)$. It is easily seen [7, Remark 1.2] that for the $n \times n$ random matrix A whose entries are i.i.d. random variables taking the values 0 with probability $1 - (2\epsilon/n)$ and $\pm\sqrt{n/(2\epsilon)}$ with probability ϵ/n each, with probability at least $1 - 2\exp(-c'\epsilon n)$,

$$\|\widetilde{A}\| \gtrsim \sqrt{\frac{n}{\epsilon}}$$

for every matrix $\widetilde{A}$ obtained by modifying A on a $c\epsilon n \times c\epsilon n$ matrix. Here, $c, c' > 0$ are absolute constants. This example shows that the dependence of the bound on $\|\widetilde{A}\|$ in Theorem 1.1 is optimal up to a possible factor of $\log \epsilon^{-1}$. In [7, Section 11], Rebrova and Vershynin asked whether this factor of $\log \epsilon^{-1}$ is necessary. We show that it is not, thereby obtaining a result which is optimal up to constants.

Theorem 1.2. *Consider an $n \times n$ random matrix A with i.i.d. entries that have zero mean and unit variance. There exist absolute constants $C_{1.2}, c_{1.2} > 0$ such that for any $\epsilon \in (0, 1/2]$, with probability at least $1 - 2\exp(-c_{1.2}\epsilon n)$, there exists an $\epsilon n \times \epsilon n$ sub-matrix of A such that replacing all of its entries with zero gives a matrix $\widetilde{A}$ with*

$$\|\widetilde{A}\| \leq C_{1.2} \sqrt{\frac{n}{\epsilon}}.$$

1.2. **Norm regularization of random symmetric matrices.** The regularization results in [2, 4] were proved for adjacency matrices of random graphs, whereas the main result of [7] holds only for random matrices with i.i.d. entries. Answering a question in [7, Section 11], our next result provides a symmetric counterpart of Theorem 1.2.

Theorem 1.3. *Consider an $n \times n$ random symmetric matrix A with i.i.d. entries on and above the diagonal that have zero mean and unit variance. There exist absolute constants $C_{1.3}, c_{1.3} > 0$ such that for any $\epsilon \in (0, 1/2]$, with probability at least $1 - 2\exp(-c_{1.3}\epsilon n)$, there exists an $\epsilon n \times \epsilon n$ sub-matrix of A such that replacing all of its entries with zero gives a matrix $\widetilde{A}$ with*

$$\|\widetilde{A}\| \leq C_{1.3} \sqrt{\frac{n}{\epsilon}}.$$

Remark. The symmetric version of the example given above shows that the dependence of the bound on $\|\widetilde{A}\|$ is optimal in terms of ϵ. Also, one can allow for the diagonals to be arbitrary independent random variables with zero mean and unit variance; this is a straightforward modification of the proof and we leave the details to the interested reader.

1.3. **Constructive norm regularization.** The norm regularization result of Rebrova and Vershynin (Theorem 1.1) is only an existential result and does not provide a way to efficiently find an appropriate $\epsilon n \times \epsilon n$ sub-matrix to zero out. In contrast, the regularization procedures of [2, 4] are algorithmic in nature. In [7, Section 11], Rebrova and Vershynin asked whether one can obtain an explicit description of an $\epsilon n \times \epsilon n$ matrix whose removal regularizes the norm.

This question was the focus of the work of Rebrova [6] who showed [6, Corollary 1.3] that for an $n \times n$ random matrix A with i.i.d. entries having a symmetric distribution and unit variance, for

any $\epsilon \in (0, 1/2]$, and for any $r \geq 1$, there is a deterministic, polynomial time algorithm to zero out an $\epsilon n \times \epsilon n$ sub-matrix in order to obtain $\widetilde{A}$ satisfying

$$\|\widetilde{A}\| \lesssim r^{3/2} \cdot \sqrt{\log \log n \cdot \log \epsilon^{-1}} \cdot \sqrt{\frac{n}{\epsilon}}$$

such that the algorithm succeeds with probability $1 - n^{0.1-r}$ (in the choice of A). Compared to the existential regularization results for i.i.d. matrices discussed earlier, this result requires that the common distribution of the entries is symmetric (as opposed to only mean 0), loses an additional factor of $\sqrt{\log \log n \cdot \log \epsilon^{-1}}$ in the bound on $\|\widetilde{A}\|$, and moreover, the failure probability (over the choice of A) of the regularization procedure is much larger than for the existential results.

Our final result remedies these shortcomings, thereby providing constructive versions of Theorems 1.2 and 1.3

Theorem 1.4. *The $\epsilon n \times \epsilon n$ matrices guaranteed in Theorems 1.2 and 1.3 can be found via a deterministic polynomial time algorithm which is guaranteed to succeed with probability at least $1 - 2\exp(-c_{1.4}\epsilon n)$ (in the choice of A). Here, $c_{1.4} > 0$ is an absolute constant.*

1.4. **Organization.** As in [7], we follow the natural high-level strategy of decomposing A into a number of parts based on the magnitude of the entries as well as controlling the operator norm of the part with the smallest entries using a version of the Grothendieck-Pietsch factorization theorem. Section 4 contains our treatment of "medium entries" and "large entries" while Section 2 and Section 3.1 contains our treatment of "small entries". Finally, these results are combined in Section 5 to prove Theorems 1.2 to 1.4.

1.5. **Notation.** Given an $n \times n$ matrix A and a subset $J \subseteq [n]$, we let A_J denote the matrix obtained from A by zeroing out the columns J^c. Given an $n \times n$ matrix A and a subset $S \subseteq [n] \times [n]$, we denote by A_S the matrix obtained from A by zeroing out all entries not in S.

For the sake of uniformity in our arguments, we will deduce all statements for $n \times n$ random matrices with i.i.d. entries above the diagonal and 0 on and below the diagonal. For brevity, we will refer to such matrices as i.i.d. random upper triangular matrices and will typically denote them by T.

1.6. **Acknowledgements.** We thank Liza Rebrova for useful discussions.

2. Controlling $2 \to \infty$ norm

As in [7], we begin by regularizing $\|\cdot\|_{2\to\infty}$ for the "small" part of the matrix.

Proposition 2.1. *Let $\epsilon \in (0, 1/2)$. Consider an $n \times n$ upper triangular matrix T with i.i.d. entries satisfying $\mathbb{E}T_{ij}^2 \leq 1$ and $|T_{ij}| \leq \sqrt{n/\log \epsilon^{-1}}$ almost surely. There exist absolute constants $C_{2.1}, c_{2.1} > 0$ such that with probability at least $1 - 2\exp(-c_{2.1}\epsilon n)$, there is a subset $J \subseteq [n]$ with $|J| \leq C_{2.1}\epsilon n$ for which*

$$\|T_{J^c}\|_{2\to\infty} \leq C_{2.1}\sqrt{n}.$$

Furthermore, this set can be found algorithmically without access to the underlying distribution.

Compared to [7, Lemma 5.1], the above proposition is valid for i.i.d. upper triangular matrices (as opposed to only i.i.d. matrices) and does not lose a factor of $\sqrt{\log \epsilon^{-1}}$ in the bound on $\|\cdot\|_{2\to\infty}$, although this comes at the cost of restricting the bound on $|T_{ij}|$ by a corresponding factor of $1/\sqrt{\log \epsilon^{-1}}$.

The key step in the proof of Proposition 2.1 is the following, which is closely related to [7, Theorem 4.2].

Lemma 2.2. *Let $\epsilon \in (0, 1/2)$ and let $\xi \in [0, n/\log \epsilon^{-1}]$ be such that $\mathbb{E}\xi \leq 1$. Let $X_1, \ldots, X_n$ as well as $Y_1, \ldots, Y_{n^2/8}$ be independent samples of ξ. There exist absolute constants $C_{2.2}, c_{2.2} > 0$ for which the following holds. For all $n \geq C_{2.2} \log \epsilon^{-1}$, there exist random variables $W_1, \ldots, W_n \in [0, 1]$ depending only on $X_1, \ldots, X_n, Y_1, \ldots, Y_{n^2/8}$ (and not the distribution of ξ) such that with probability at least $1 - 2\exp(-c_{2.2} n \log \epsilon^{-1})$ over $Y_1, \ldots, Y_{n^2/8}$, we have*

$$\sum_{j=1}^{n} W_j X_j \leq C_{2.2} n \quad \text{almost surely over } X_1, \ldots, X_n, \text{ and}$$

$$1 \leq \mathbb{E}_{X_1,\ldots,X_n}\left(\prod_{i=1}^{n} W_i\right)^{-1} \leq 1 + \epsilon.$$

Proof. Let $K = 1/\log \epsilon^{-1}$ so that $\xi \in [0, Kn]$ and $n = \Omega(1/K)$. Let $t = \lceil \log_2(Kn) \rceil$. Let $q_{-1} = q_0 = 0, q_{t+1} = Kn$, and for $k \in [t]$, let

$$q_k = \sup\{r \geq 0 : \mathbb{P}[\xi \geq r] \geq 2^{-k}\}.$$

Let $\widehat{\mathbb{P}}$ denote the empirical measure generated by $Y_1, \ldots, Y_{n^2/8}$. Let $\widehat{q}_0 = 0, \widehat{q}_{t+1} = 2Kn$, and for $k \in [t]$, let

$$\widehat{q}_k = \sup\{r \geq 0 : \widehat{\mathbb{P}}[[r, \infty)] \geq 2^{-k}\}.$$

Let $\mathcal{G}$ denote the event, measurable with respect to $Y_1, \ldots, Y_{n^2/8}$, that

$$q_{k-1} \leq \widehat{q}_k \leq q_{k+1} \quad \forall k \in [t].$$

Then, it follows from a staightforward application of the Chernoff bound and the union bound that

$$\mathbb{P}_{Y_1,\ldots,Y_{n^2/8}}[\mathcal{G}] \geq 1 - \sum_{k=1}^{t} \exp\left(-\Omega\left(\frac{n^2}{2^k}\right)\right) \geq 1 - 2\exp(-\Omega(n/K)).$$

Now, fix a realisation of $Y_1, \ldots, Y_{n^2/8}$ and the associated quantities $\widehat{q}_0, \ldots, \widehat{q}_{t+1}$. Let $k_0 = 0$ and inductively define

$$k_{i+1} = \max\{k_i + 1, \max\{j \in (k_i, t+1] : \widehat{q}_j < 2\widehat{q}_{k_i}\}\}.$$

Here, the maximum of the empty set is $-\infty$. Let ℓ denote the first index for which $k_\ell = t+1$. Note that for all $i \in \{0, 1, \ldots, \ell - 1\}$,

$$2^{-k_i}\widehat{q}_{k_{i+1}} \lesssim 2^{-k_i}\widehat{q}_{k_i} + 2^{-k_{i+1}}\widehat{q}_{k_{i+1}}.$$

Let $L \geq 1$ be a sufficiently large absolute constant to be specified later. For a realisation of $X_1, \ldots, X_n$ and for $i \in \{0, 1, \ldots, \ell - 1\}$, let

$$\nu_i = \#\{j \in [n] : X_j \in [\widehat{q}_{k_i}, \widehat{q}_{k_{i+1}})\} \text{ and } w_i = \min\left(1, \frac{L2^{-k_i} n}{\nu_i}\right).$$

Also, for $j \in [n]$, we let $i(j)$ denote the unique (by construction) index $i \in \{0, 1, \ldots \ell - 1\}$ for which $X_j \in [\widehat{q}_{k_i}, \widehat{q}_{k_{i+1}})$. For $j \in [n]$ and $i \in \{0, 1, \ldots, \ell - 1\}$, let

$$W_j = w_{i(j)} \text{ and } Z_i = \prod_{j:i(j)=i} W_j^{-1}.$$

Then, on the event $\mathcal{G}$, we have

$$\sum_{j=1}^{n} W_j X_j = \sum_{i=0}^{\ell-1} \sum_{j:i(j)=i} W_j X_j \leq \sum_{i=0}^{\ell-1} Ln \cdot 2^{-k_i}\widehat{q}_{k_{i+1}} \lesssim \sum_{k=0}^{t+1} Ln \cdot 2^{-k}\widehat{q}_k$$

4

$$\lesssim Ln + \sum_{k=1}^{t} Ln \cdot 2^{-k} q_{k+1} \lesssim Ln \left(1 + \mathbb{E}\xi\right) \lesssim Ln.$$

Moreover, on the event $\mathcal{G}$, we have for any $i \in \{0, 1, \ldots, \ell - 1\}$ that

$$\mathbb{E}_{X_1, \ldots, X_n}[Z_i] \leq 1 + \mathbb{E}\left[\left(\frac{\nu_i}{L 2^{-k_i} n} \right)^{\nu_i} \mathbb{1}_{\nu_i > L 2^{-k_i} n} \right]$$

$$\leq 1 + \sum_{s > L 2^{-k_i} n} \left(\frac{s}{L 2^{-k_i} n} \right)^s \mathbb{P}[\nu_i = s]$$

$$\leq 1 + \sum_{s > L 2^{-k_i} n} \left(\frac{s}{L 2^{-k_i} n} \right)^s \binom{n}{s} \mathbb{P}[\xi \geq \widehat{q}_{k_i}]^s$$

$$\leq 1 + \sum_{s > L 2^{-k_i} n} \left(\frac{e}{L 2^{-k_i}} \right)^s \mathbb{P}[\xi \geq q_{k_i - 1}]^s$$

$$\leq 1 + \sum_{s > L 2^{-k_i} n} \left(\frac{4e}{L} \right)^s \leq \exp\left(2 \left(\frac{8e}{L} \right)^{\frac{Ln}{2^{2+k_i}}} \right),$$

for a sufficiently large absolute constant L. Then, since the random variables $|\nu_0|, \ldots, |\nu_{\ell-1}|$ are negatively associated and since Z_i is an increasing function of $|\nu_i|$, it follows that

$$\mathbb{E}_{X_1, \ldots, X_n} \left[\prod_{j=1}^{n} W_j^{-1} \right] = \mathbb{E}\left[\prod_{i=0}^{\ell-1} Z_i \right] \leq \prod_{i=0}^{\ell-1} \mathbb{E}[Z_i]$$

$$\leq \exp\left(2 \left(\frac{8e}{L} \right)^{L/(8K)} \right) \leq 1 + \epsilon$$

for a sufficiently large absolute constant L, where the final inequality uses that $K = 1/\log \epsilon^{-1}$. $\quad\square$

We can now quickly deduce Proposition 2.1.

Proof of Proposition 2.1. We may assume that $n = \Omega(\epsilon^{-1})$ as otherwise, the desired probability bound is negative. Moreover, by adding an extra row and column of zeros (if necessary), we may assume that n is even. Recall that the diagonal entries of T are 0. Consider the $n/2 \times n/2$ matrices

$$N_{ij} = T_{i, n/2 + j}, \text{ and}$$

$$N'_{ij} = \begin{cases} T_{i + n/2, j + n/2} & \text{if } j > i \text{ and } i < n/2 \\ T_{n/2 - i, n/2 - j + 1} & \text{if } j \leq i \text{ and } i < n/2 \\ 0 & \text{if } i = n/2. \end{cases}$$

Then, it is straightforward to see that if we can find $O(\epsilon n)$ columns in N and N' to zero-out such that the resulting matrices have $\|\cdot\|_{2 \to \infty}$ norm $O(\sqrt{n})$, then the same is true for T and in fact, the choice of columns used for N and N' correspond to an obvious choice of columns for T.

By taking the union bound, it suffices to show the following: with probability at least $1 - 2\exp(-\Omega(n \log \epsilon^{-1}))$ over the realisation of N', with probability at least $1 - 2\exp(-\Omega(\epsilon n))$ over the realization of N, we can algorithmically find a set of $O(\epsilon n)$ columns of N to zero-out such that the resulting matrix has $\|\cdot\|_{2 \to \infty}$ norm $O(\sqrt{n})$. This follows from a direct application of Lemma 2.2.

Indeed, we treat the entries $(N'_{ij})^2$ for $i < n/2$ as the i.i.d. samples $Y_1, \ldots, Y_{n^2/8}$. Then, by Lemma 2.2, with probability at least $1 - 2\exp(-\Omega(n \log \epsilon^{-1}))$ over the realization of N', the following

holds. For any $i \in [n/2]$, we can find $W_{i1}, \ldots, W_{i,n/2} \in [0,1]$ depending only on $N_{i1}, \ldots, N_{i,n/2}$ and N' such that

$$\sum_{j=1}^{n/2} W_{ij} N_{ij}^2 \leq C_{2.2} n \text{ and}$$

$$1 \leq \mathbb{E}_{N_{i1}, \ldots, N_{i,n/2}} \left(\prod_{j=1}^{n/2} W_{ij} \right)^{-1} \leq \exp(\epsilon).$$

Let

$$V_j = \prod_{i=1}^{n/2} W_{ij}, \quad j \in [n/2].$$

Then, $W_{ij} \leq V_j$ for all $i \in [n/2]$ so that

$$\sum_{j=1}^{n/2} V_j N_{ij}^2 \leq C_{2.2} n \quad \forall i \in [n/2] \text{ and} \tag{2.1}$$

$$1 \leq \mathbb{E}_N \left(\prod_{j=1}^{n/2} V_j \right)^{-1} \leq \exp(\epsilon n). \tag{2.2}$$

Let $J = \{ j \in [n/2] : V_j < e^{-2} \}$. By (2.2) and Markov's inequality, it follows that $|J| \leq \epsilon n$ with probability at least $1 - \exp(-\epsilon n)$ over the choice of N. Moreover, by (2.1), it follows that for all $i \in [n/2]$, $\sum_{j \in J^c} N_{ij}^2 \leq e^2 \cdot C_{2.2} n$, which completes the proof. $\qquad\square$

3. Controlling $\infty \to 2$ and operator norms

3.1. $\infty \to 2$ **norm.** In this subsection, we will show how to turn a bound on $\|\cdot\|_{2 \to \infty}$ into a bound on $\|\cdot\|_{\infty \to 2}$. In contrast to the corresponding step in [7], we will be able to accomplish this *without* removing any additional columns, which will be useful for our algorithmic regularization procedure. The goal of this subsection is to prove the following Seginer-type [10] result for the $\|\cdot\|_{\infty \to 2}$ norm.

Proposition 3.1. *Let T be an $n \times n$ upper triangular matrix with i.i.d. entries of mean 0 and variance at most 1. There exists an absolute constant $C_{3.1} > 0$ such that with probability at least $1 - 4^{-n}$, the following holds:*

$$\|T_I\|_{\infty \to 2} \leq C_{3.1}(\sqrt{n} \|T_I\|_{2 \to \infty} + n) \text{ for all } I \subseteq [n], |I| \geq n/2.$$

The proof of Proposition 3.1 will be presented at the end of this subsection following a series of preparatory lemmas. We begin with the following tight relationship between $\|\cdot\|_{2 \to \infty}, \|\cdot\|_{\infty \to 2}$, and $\|\cdot \mathbf{1}\|_2$, proved in [7] for a completely i.i.d. matrix.

Lemma 3.2 ([7, Lemmas 6.3 and 6.4]). *Let A be a random $n \times n$ matrix with i.i.d. entries. Then*

$$\mathbb{E}\|A\|_{\infty \to 2} \leq C_{3.2}(\sqrt{n}\mathbb{E}\|A\|_{2 \to \infty} + \mathbb{E}\|A\mathbf{1}\|_2),$$

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

$$\|A\|_{\infty \to 2} \leq C_{3.2}(\sqrt{n}\mathbb{E}\|A\|_{2 \to \infty} + \mathbb{E}\|A\mathbf{1}\|_2).$$

We establish an analogous version for random upper triangular matrices with i.i.d. entries.

Lemma 3.3. *Let T be a random upper triangular $n \times n$ matrix with i.i.d. entries, and let J be a subset of columns with $|J| \geq n/2$. Then with probability at least $1 - 2e^{-n}$,*

$$\|T_J\|_{\infty \to 2} \leq C_{3.3}(\sqrt{n}\mathbb{E}\|T_J\|_{2 \to \infty} + \mathbb{E}\|T_J\mathbf{1}\|_2).$$

During the course of our proof, we will make use of the following estimates regarding random matrix models of different "shapes".

Lemma 3.4. *Let $R \subseteq S \subseteq [n] \times [n]$. Then*

$$\mathbb{E}\|A_R\|_{\infty \to 2} \leq 2\mathbb{E}\|A_S\|_{\infty \to 2}$$

and

$$\mathbb{E}\|A_R \mathbf{1}\|_2 \leq 2\mathbb{E}\|A_S \mathbf{1}\|_2.$$

Proof. Let $\mu = \mathbb{E}\xi$. We have

$$
\begin{aligned}
\mathbb{E}\|A_R\|_{\infty \to 2} &= \mathbb{E}\|A_R + \mathbb{E}[A_{S \setminus R} - \mathbb{E}A_{S \setminus R}]\|_{\infty \to 2} \\
&\leq \mathbb{E}\|A_R + A_{S \setminus R} - \mathbb{E}A_{S \setminus R}\|_{\infty \to 2} \leq \mathbb{E}\|A_S\|_{\infty \to 2} + \|\mu \mathbf{1}_{S \setminus R}\|_{\infty \to 2} \\
&\leq \mathbb{E}\|A_S\|_{\infty \to 2} + \|\mu \mathbf{1}_S\|_{\infty \to 2} = \mathbb{E}\|A_S\|_{\infty \to 2} + \|\mathbb{E}A_S\|_{\infty \to 2} \\
&\leq 2\mathbb{E}\|A_S\|_{\infty \to 2},
\end{aligned}
$$

where we have used Jensen's inequality twice. An analogous proof establishes the second inequality as well. □

We can now prove Lemma 3.3.

Proof of Lemma 3.3. We may write $T_J = A_S$ for some $S \subseteq [n] \times [n]$, where note that S contains an $\lfloor n/4 \rfloor \times \lfloor n/4 \rfloor$ (not necessarily consecutive) block as a subset. Since zeroing out entries of a matrix cannot increase the $\|\cdot\|_{2 \to \infty}$ norm, it follows that

$$\mathbb{E}\|A\|_{2 \to \infty} \lesssim \mathbb{E}\|A_{\lfloor n/4 \rfloor \times \lfloor n/4 \rfloor}\|_{2 \to \infty} \leq \mathbb{E}\|T_J\|_{2 \to \infty} \leq \mathbb{E}\|A\|_{2 \to \infty},$$

where the first inequality uses the triangle inequality along with the previously mentioned fact about zeroing out entries. Moreover, we have

$$\mathbb{E}\|A\|_{\infty \to 2} \lesssim \mathbb{E}\|A_{\lfloor n/4 \rfloor \times \lfloor n/4 \rfloor}\|_{\infty \to 2} \lesssim \mathbb{E}\|T_J\|_{\infty \to 2} \lesssim \mathbb{E}\|A\|_{\infty \to 2},$$

where the first inequality uses the triangle inequality and Lemma 3.4 and the subsequent inequalities use Lemma 3.4. Similarly,

$$\mathbb{E}\|A\mathbf{1}\|_2 \lesssim \mathbb{E}\|A_{\lfloor n/4 \rfloor \times \lfloor n/4 \rfloor}\mathbf{1}\|_2 \lesssim \mathbb{E}\|T_J\mathbf{1}\|_2 \lesssim \mathbb{E}\|A\mathbf{1}\|_2.$$

Next, write $A = T_J + A_{S'}$ for $S' = ([n] \times [n]) \setminus S$. Then, for $t \geq 2\mathbb{E}\|A_{S'}\|_{\infty \to 2}$, we have

$$
\begin{aligned}
\mathbb{P}[\|A\|_{\infty \to 2} \geq t] &\geq \mathbb{P}[\|A_{S'}\|_{\infty \to 2} \leq t] \cdot \mathbb{P}[\|T_J\|_{\infty \to 2} \geq 2t] \\
&\geq \mathbb{P}[\|T_J\|_{\infty \to 2} \geq 2t]/2,
\end{aligned}
$$

where the second line follows from Markov's inequality. Since

$$\mathbb{E}\|A_{S'}\|_{\infty \to 2} \lesssim \mathbb{E}\|A\|_{\infty \to 2} \lesssim \sqrt{n}\mathbb{E}\|A\|_{2 \to \infty} + \mathbb{E}\|A\mathbf{1}\|_2 \lesssim \sqrt{n}\mathbb{E}\|T_J\|_{2 \to \infty} + \mathbb{E}\|T_J\mathbf{1}\|_2,$$

where the first inequality is by Lemma 3.4, the second inequality is by Lemma 3.2, and the third is by the previously established inequalities, we may choose

$$t = C(\sqrt{n}\mathbb{E}\|T_J\|_{2 \to \infty} + \mathbb{E}\|T_J\mathbf{1}\|_2)$$

with C is a large absolute constant guaranteeing that

$$t \geq (2 + C_{3.2}) \max(\sqrt{n}\mathbb{E}\|A\|_{2 \to \infty} + \mathbb{E}\|A\mathbf{1}\|_2, \mathbb{E}\|A_{S'}\|_{\infty \to 2}).$$

Finally, for such a choice of C and t, we have by Lemma 3.2 that

$$\mathbb{P}[\|T_J\|_{\infty \to 2} \geq 2t] \leq 2\mathbb{P}[\|A\|_{\infty \to 2} \geq t] \leq 2e^{-n},$$

as desired. □

With the preceding lemma in hand, we can prove the following, which shows that with at least some exponentially *small* probability, the $\|\cdot\|_{2\to\infty}$ and $\|\cdot\|_{\infty\to2}$ norms of the matrix are already regularized.

Lemma 3.5. *Let T be a random upper triangular $n \times n$ matrix with i.i.d. entries of mean 0 and variance at most 1, and let J be a subset of columns with $|J| \geq n/2$. For any $\delta \in (0, 1/2)$,*

$$\|T_J\|_{2\to\infty} \leq C_{3.5}\delta^{-1}\sqrt{n} \quad \text{and} \quad \|T_J\|_{\infty\to2} \leq C_{3.5}\delta^{-1}n$$

with probability at least $\exp(-\delta^2 n)/2$.

Proof. The proof is identical to that of [7, Lemma 6.5] with the application of [7, Lemma 6.4] replaced by Lemma 3.3. □

Finally, we need the following symmetrization estimate from [7].

Lemma 3.6 (Proof of [7, Lemma 6.1]). *Let A be an $n \times n$ matrix and let $\widetilde{A}$ denote the random matrix with entries $\widetilde{A}_{ij} = \epsilon_{ij}A_{ij}$, where ϵ_{ij} are i.i.d. Rademacher random variables. There exists an absolute constant $c_{3.6} > 0$ such that for all $t \geq 1$, with probability at least $1 - 2^n \exp(-c_{3.6}tn)$,*

$$\|\widetilde{A}\|_{\infty\to2} \leq \sqrt{2tn}\|A\|_{2\to\infty}.$$

We now have all the ingredients needed to prove Proposition 3.1.

Proof of Proposition 3.1. Fix $I \subseteq [n]$ with $|I| \geq n/2$, let $B = T_I$, and let $\widetilde{B} = B - B'$, where B' denotes an independent copy of T_I. For $t \geq 1$ and $\delta \in (0, 1/2)$, let

$$\mathcal{E}(t) := \{\|\widetilde{B}\|_{\infty\to2} \leq \sqrt{2tn}\|\widetilde{B}\|_{2\to\infty}\},$$
$$\mathcal{F}(\delta) := \{\|B'\|_{2\to\infty} \leq C_{3.5}\delta^{-1}\sqrt{n} \wedge \|B'\|_{\infty\to2} \leq C_{3.5}\delta^{-1}n\}, \quad \text{and}$$
$$\mathcal{G}(t,\delta) := \{\|B\|_{\infty\to2} \leq 2\sqrt{2tn}\|B\|_{2\to\infty} + (2C_{3.5}\delta^{-1}\sqrt{2t} + C_{3.5}\delta^{-1})n\}.$$

Then, by the triangle inequality, we have that on the event $\mathcal{G}(t,\delta)^c \cap \mathcal{F}(\delta)$,

$$\|\widetilde{B}\|_{2\to\infty} \leq \|B\|_{2\to\infty} + \|B'\|_{2\to\infty} \leq \|B\|_{2\to\infty} + C_{3.5}\delta^{-1}\sqrt{n}, \quad \text{and}$$
$$\|\widetilde{B}\|_{\infty\to2} \geq \|B\|_{\infty\to2} - \|B'\|_{\infty\to2} \geq 2\sqrt{2tn}\left(\|B\|_{2\to\infty} + C_{3.5}\delta^{-1}\sqrt{n}\right).$$

In particular, this implies

$$\mathcal{G}(t,\delta)^c \cap \mathcal{F}(\delta) \subseteq \mathcal{E}^c(t).$$

Therefore, since $\mathcal{G}(t,\delta)$ and $\mathcal{F}(\delta)$ are independent, we have

$$\mathbb{P}[\mathcal{G}(t,\delta)^c] \leq \frac{\mathbb{P}[\mathcal{E}^c(t)]}{\mathbb{P}[\mathcal{F}(\delta)]} \leq \frac{2^n \exp(-c_{3.6}tn)}{\exp(-\delta^2 n)/2},$$

where we have used Lemmas 3.5 and 3.6. Finally, choosing $t = 1000/c_{3.6}$, $\delta = 1/4$, and taking the union bound over at most 2^n choices of I, we obtain the desired conclusion. □

3.2. Operator norm.
Having thus established control over $\|\cdot\|_{\infty\to2}$, we can apply a version of the Grothendieck-Pietsch theorem [5, Proposition 15.11] as in [4,7] to establish control over the operator norm.

Theorem 3.7. *Let B be a $k \times m$ real matrix and let $\delta > 0$. There exists $J \subseteq [m]$ with $|J| \leq \delta m$ for which*

$$\|B_{J^c}\| \leq \frac{2}{\sqrt{\delta m}}\|B\|_{\infty\to2}.$$

Now, applying Propositions 2.1 and 3.1 and Theorem 3.7 in sequence immediately yields the following result.

Proposition 3.8. *Let $\epsilon \in (0, 1/2)$ and consider an $n \times n$ upper triangular matrix T with i.i.d. mean 0 entries of variance at most 1, and such that $|T_{ij}| \leq \sqrt{n/\log(1/\epsilon)}$ almost surely. Then with probability at least $1 - 2\exp(-c_{3.8}\epsilon n)$, there is a subset $J \subseteq [n]$ with $|J| \leq C_{3.8}\epsilon n$ such that*

$$\|T_{J^c}\| \leq C_{3.8}\sqrt{\frac{n}{\epsilon}}.$$

4. Medium and Large Entries

In this section, we handle the "medium" and "large" entries. Let $\epsilon \in (0, 1/2)$, and recall that Proposition 3.8 takes care of those entries of the matrix T which have absolute value at most $\sqrt{n/\log \epsilon^{-1}}$. We split the remaining entries of T into three separate classes:

- $|T_{i,j}| \in \left(\sqrt{n/\log \epsilon^{-1}}, \sqrt{n/(\epsilon \log^2 \epsilon^{-1})}\right]$,
- $|T_{i,j}| \in \left(\sqrt{n/(\epsilon \log^2 \epsilon^{-1})}, 5\sqrt{n/\epsilon}\right]$, and
- $|T_{i,j}| > 5\sqrt{n/\epsilon}$.

The first and third classes are handled using arguments similar to [7]. The second class requires a more intricate argument, and will be considered at the end of this section.

We will need an elementary lemma from [7] which will allow us to combine different sub-matrices that we zero out into one.

Lemma 4.1 ([7, Lemma 8.7]). *Let M be an $n \times n$ matrix. Let $I, J \subseteq [n]$. Then*

$$\|B - B_{I \times J}\| \leq \|B_{I^c \times [n]}\| + \|B_{[n] \times J^c}\| \leq 2\|B\|.$$

Proof. Indeed,

$$\|B - B_{I \times J}\| \leq \|B_{I^c \times [n]}\| + \|B_{I \times J^c}\| \leq \|B_{I^c \times [n]}\| + \|B_{[n] \times J^c}\| \leq 2\|B\|,$$

where the first inequality is the triangle inequality, and the remaining inequalities use that zeroing out a subset of rows and/or columns cannot increase the operator norm. $\square$

We also require the following bound on the operator norm of the matrix in terms of the ℓ_1 norms of its rows and columns, due to Schur [9].

Lemma 4.2. *For any matrix A,*

$$\|A\| \leq \|A\|_{\infty \to 1}^{1/2} \|A^T\|_{\infty \to 1}^{1/2}.$$

Now, we proceed to the treatment of the three classes of entries. The easiest is the third class, for which the result follows by a straightforward application of the Chernoff bound.

Lemma 4.3 (Modification of [7, Corollary 8.6]). *Let $\epsilon \in (0, 1/2]$. Consider an $n \times n$ random matrix upper triangular matrix T with independent (but not necessarily identically distributed) entries T_{ij} such that $T_{ij} = 0$ or $|T_{ij}| \geq 5\sqrt{n/\epsilon}$, and which satisfy $\mathbb{E}T_{ij}^2 \leq 1$. Then with probability $1 - \exp(-\epsilon n)$, all nonzero entries of T are contained in an $\epsilon n \times \epsilon n$ matrix.*

Remark. In [7], the above lemma is stated for identically distributed entries but as can be seen from the proof, this assumption is unnecessary.

Next, we handle entries in the first class. The following lemma, and its proof (which we include for completeness), are essentially identical to [7, Proposition 8.4].

9

Lemma 4.4. *Let $\epsilon \in (0, 1/2]$. Let T be a random upper triangular $n \times n$ matrix with i.i.d. entries of mean 0 and variance at most 1 such that $T_{ij} = 0$ or $|T_{ij}| \in [\sqrt{n/\log \epsilon^{-1}}, \sqrt{n/(\epsilon \log^2 \epsilon^{-1})}]$ almost surely. With probability at least $1 - 2\exp(-\epsilon n/4)$ we can zero out an $\epsilon n \times \epsilon n$ block to obtain $\widetilde{T}$ satisfying*

$$\|\widetilde{T}\| \leq C_{4.4}\sqrt{\frac{n}{\epsilon}}.$$

Proof. Let $B_{ij} = \mathbb{1}_{T_{ij} \neq 0}$. Since the B_{ij} are independent Bernoulli random variables with mean $p \leq (\log \epsilon^{-1})/n$, it follows from [7, Corollary 8.2] that with probability at least $1 - 2\exp(-\epsilon n/4)$, we can find an $\epsilon n \times \epsilon n$ sub-matrix of B to zero out so that the rows and columns of the resulting matrix have at most $O(\log \epsilon^{-1})$ ones each. Since $|T_{ij}| \leq \sqrt{n/(\epsilon \log^2 \epsilon^{-1})}$, this shows that we can find an $\epsilon n \times \epsilon n$ sub-matrix of T to zero out so that the rows and columns of the resulting matrix have ℓ_1 norm at most $O(\sqrt{n/\epsilon})$ each. An application of Lemma 4.2 gives the desired result. $\square$

Finally, we handle entries in the second class.

Proposition 4.5. *Let $\epsilon \in (0, 1/2]$. Let T be a random upper triangular $n \times n$ matrix with i.i.d. entries of mean 0 and variance at most 1 such that $T_{ij} = 0$ or $|T_{ij}| \in [\sqrt{n/(\epsilon \log^2 \epsilon^{-1})}, 5\sqrt{n/\epsilon}]$ almost surely. With probability at least $1 - 2\exp(-c_{4.5}\epsilon n)$ we can zero out a square sub-matrix of size $C_{4.5}\epsilon n \times C_{4.5}\epsilon n$ to obtain $\widetilde{T}$ satisfying*

$$\|\widetilde{T}\| \leq C_{4.5}\sqrt{\frac{n}{\epsilon}}.$$

Proof. As always, we can assume $n = \Omega(\epsilon^{-1})$ since otherwise, the target probability is negative. Consider the matrix $B_{ij} = \mathbb{1}_{T_{ij} \neq 0}$. This is an upper triangular matrix with i.i.d. Ber(p) entries, where $p \leq (\epsilon \log^2 \epsilon^{-1})/n$. We will show that with probability $1 - 2\exp(-c\epsilon n)$, one can remove $O(\epsilon n)$ rows of B to obtain B' such that each row and each column of B' has at most a single entry equal to 1. Since $|T_{ij}| \leq 5\sqrt{n/\epsilon}$, this shows that the corresponding matrix T' will satisfy $\|T'\| \leq 5\sqrt{n/\epsilon}$. By using the symmetry $(i,j) \mapsto (n-j, n-i)$, it follows that with probability $1 - 2\exp(-c\epsilon n)$, one can remove $O(\epsilon n)$ columns of T to obtain T'' with $\|T''\| \leq 5\sqrt{n/\epsilon}$. Finally, Lemma 4.1 shows that the matrix $\widetilde{T}$ obtained by removing the intersection of these $O(\epsilon n)$ rows and $O(\epsilon n)$ columns will have operator norm bounded by $O(\sqrt{n/\epsilon})$ as desired.

Therefore, consider B as above. Let t denote the number of rows of B with exactly one entry equal to 1 and let b denote the number of rows of B with at least two entries equal to 1. Also, let $B^{(1)}$ denote the $t \times n$ sub-matrix of B in which each row has exactly one entry equal to 1. Then, by the Chernoff–Hoeffding bound [3, Theorem 1], except with probability at most $2\exp(-c\epsilon n \log^2 \epsilon^{-1})$, we have $t \leq t^* := \lceil 2\epsilon n \log^2 \epsilon^{-1} \rceil$ and $b = O(\epsilon n)$. Hence, with probability at least $1 - 2\exp(-c\epsilon n \log^2 \epsilon^{-1})$, we have enough room to remove all rows of B with at least two entries equal to 1.

It remains to deal with $B^{(1)}$. Let $B^{(-)}$ denote the sub-matrix of $B^{(1)}$ consisting of the first n_- columns and $B^{(+)}$ denote the sub-matrix of $B^{(1)}$ consisting of the remaining n_+ columns, where n_- and n_+ are chosen to be as close to each other as possible. We wish to bound the probability

$$\mathbb{P}[B^{(\pm)} \in \mathcal{B} \mid t]$$

uniformly for all $t \leq t^*$, where $\mathcal{B}$ denotes the collection of $\{0,1\}$-valued $t \times n_\pm$ matrices for which the number of non-zero entries present in the union of all columns with at least two non-zero entries is not $O(\epsilon n)$. We present the bound for $B^{(-)}$, noting that the same argument also applies to $B^{(+)}$.

10

Clearly, the probability is maximized for $t = t^*$. Moreover, since the rows of $B^{(-)}$ are independent and since each row has the distribution

$$R \sim \begin{cases} \mathbf{0}, & \text{w.p. } n_-/n, \\ e_i, & \text{w.p. } n_+/(n_- \cdot n) \text{ for all } i \in [n_-], \end{cases}$$

it follows that for all $t \leq t^*$ and $n \geq 10$,

$$\mathbb{P}[B^{(-)} \in \mathcal{B} \mid t] \leq \mathbb{P}[B^{(-)} \in \mathcal{B} \mid t^*] \leq \mathbb{P}[\widehat{B} \in \mathcal{B}],$$

where $\widehat{B}$ denotes a random $t^* \times n_-$ matrix whose entries are i.i.d. $\mathrm{Ber}(2/n)$ random variables. We note that the above trick of splitting into $B^{(\pm)}$ and passing to the independent model is closely related to the proof of [1, Theorem 4.1].

To bound $\mathbb{P}[\widehat{B} \in \mathcal{B}]$, we note two things. First, by the Chernoff-Hoeffding bound, with probability at least $1 - 2\exp(-c\epsilon n)$, there are at most $s^* = \epsilon n / \log \epsilon^{-1}$ columns of $\widehat{B}$ with at least two non-zero entries. Second, assuming $s^* \geq 1$ (otherwise, we are already done) by the Chernoff-Hoeffding bound and the union bound, the probability that *any* $t^* \times \lfloor s^* \rfloor$ block of $\widehat{B}$ has more than $O(\epsilon n)$ entries is at most $2\exp(-c\epsilon n)$. Combining these two facts shows that $\mathbb{P}[\widehat{B} \in \mathcal{B}] \leq 4\exp(-c\epsilon n)$.

To summarize, we have shown that except with probability $O(\exp(-c\epsilon n))$, the following hold simultaneously: $b = O(\epsilon n)$, $B^{(-)} \notin \mathcal{B}$, $B^{(+)} \notin \mathcal{B}$. On this event, we are guaranteed that there are $O(\epsilon n)$ rows of B which either have at least two non-zero entries, or which contain a non-zero entry in a column of $B^{(\pm)}$ with at least two non-zero entries. Then, we can simply zero out all such rows to obtain the desired conclusion. $\qquad\square$

5. Proof of Theorems 1.2 to 1.4

We are now in position to prove all our results.

Proof of Theorems 1.2 to 1.4. We may assume that $n = \Omega(\epsilon^{-1})$ since otherwise, the desired success probability is negative and the statements are vacuously true. Moreover, it suffices to prove a version of Theorem 1.4 for i.i.d. random upper triangular matrices since this readily implies Theorems 1.2 to 1.4. We decompose

$$T = S + M_1 + M_2 + L,$$

where S contains the entries with magnitude at most $\sqrt{n/\log \epsilon^{-1}}$, M_1 contains the entries with magnitude in

$$I_1 = (\sqrt{n/\log \epsilon^{-1}}, \sqrt{n/(\epsilon \log^2 \epsilon^{-1})}],$$

M_2 contains the entries with magnitude in

$$I_2 = (\sqrt{n/(\epsilon \log^2 \epsilon^{-1})}, 5\sqrt{n/\epsilon}],$$

and L contains the entries with magnitude greater than $5\sqrt{n/\epsilon}$.

Since $(M_1)_{ij} = T_{ij} \mathbb{1}_{|T_{ij}| \in I_1}$, we see that M_1 satisfies the assumptions of Lemma 4.4. Therefore, with probability at least $1 - 2\exp(-\epsilon n/4)$, we can find an $\epsilon n \times \epsilon n$ sub-matrix of M_1 to zero out so that the resulting matrix $\|\widetilde{M_1}\|$ has norm at most $C_{4.4}\sqrt{n/\epsilon}$. Moreover, the proof of [7, Corollary 8.2] shows that this step can be done algorithmically. Indeed, one only needs to find the rows and columns of M_1 with more than $O(\log \epsilon^{-1})$ non-zero entries and then zero out *any* $\epsilon n \times \epsilon n$ sub-matrix containing all the entries contained in such rows and columns. This last step is easy since [7, Corollary 8.2] guarantees that there are at most ϵn such entries so that they can be trivially placed in an $\epsilon n \times \epsilon n$ sub-matrix.

Similarly, we can apply Proposition 4.5 to M_2 in order to find a suitable $\epsilon n \times \epsilon n$ sub-matrices to remove. In order to do this algorithmically, let R_1 denote the set of rows with exactly one non-zero

entry, let R_2 denote the set of rows with at least two non-zero entries, let $C_2(R_1)$ denote the set of columns with at least two non-zero entries in rows in R_1, and let R_1' denote the set of rows in R_1 such that the unique non-zero entry is in a column in $C_2(R_1)$. Define $C_1, C_2, R_2(C_1), C_1'$ similarly with the role of rows and columns interchanged. Then, the proof of Proposition 4.5 shows that, with probability at least $1 - 2\exp(-c_{4.5}\epsilon n)$, the intersection of the rows $R_2 \cup R_1'$ and the columns $C_2 \cup C_1'$ is contained in an $C_{4.5}\epsilon n \times C_{4.5}\epsilon n$ sub-matrix and zeroing it out suffices to regularize the norm of M_2.

To regularize the norm of L, we can apply Lemma 4.3. Again, this is algorithmic, since we only need to remove a sub-matrix containing all non-zero entries (of which there are at most ϵn with probability at least $1 - \exp(-\epsilon n)$, see the proof of [7, Corollary 8.6]).

Finally, we regularize the norm of S. Let $\xi' = \xi\mathbb{1}_{|\xi| \le \sqrt{n/\log\epsilon^{-1}}}$ where ξ denotes the common distribution of the upper triangular entries of T and note that the upper triangular entries of S are i.i.d. copies of ξ'. Since $\mathbb{E}[\xi'^2] \le \mathbb{E}[\xi^2] \le 1$, we can apply Proposition 2.1 to algorithmically find (with probability at least $1 - 2\exp(-c_{2.1}\epsilon n)$) a set $J_0 \subseteq [n]$ of $O(\epsilon n)$ columns such that

$$\|S_{J_0^c}\|_{2\to\infty} = O(\sqrt{n}).$$

In order to be able to apply Proposition 3.1, we need such a result not for S but for $S - \mathbb{E}S$. For this, we note that

$$|\mathbb{E}\xi'| = \left|\mathbb{E}\xi\mathbb{1}_{|\xi|\le\sqrt{n/\log\epsilon^{-1}}}\right| = \left|\mathbb{E}\left[\xi - \xi\mathbb{1}_{|\xi|>\sqrt{n/\log\epsilon^{-1}}}\right]\right|$$

$$= \left|\mathbb{E}\xi\mathbb{1}_{|\xi|>\sqrt{n/\log\epsilon^{-1}}}\right| \le (\mathbb{E}\xi^2)^{1/2}\mathbb{P}\left[|\xi| > \sqrt{n/\log\epsilon^{-1}}\right]^{1/2} \le \sqrt{\frac{\log\epsilon^{-1}}{n}}.$$

Therefore, for all $J \subseteq [n]$,

$$\|(\mathbb{E}S)_J\|_{2\to\infty} \le \sqrt{\log\epsilon^{-1}}, \quad \|(\mathbb{E}S)_J\| \le \sqrt{n\log\epsilon^{-1}}.$$

In particular,

$$\|(S - \mathbb{E}S)_{J_0^c}\|_{2\to\infty} = O(\sqrt{n} + \sqrt{\log\epsilon^{-1}}) = O(\sqrt{n}).$$

Combining this with Proposition 3.1, we see that except with probability at least $1 - 4^{-n} - 2\exp(-c_{2.1}\epsilon n)$,

$$\|(S - \mathbb{E}S)_{J_0^c}\|_{\infty\to 2} = O(n). \tag{5.1}$$

Therefore, by Theorem 3.7, we see that there exists some μ with $|\mu| \le \sqrt{\log\epsilon^{-1}/n}$ and a subset of columns J_1 with $J_1 = O(\epsilon n)$ such that for $J_* = J_1 \cup J_0$,

$$\|(S - \mu\mathbf{1}\mathbf{1}^T)_{J_*^c}\| = O(\sqrt{n/\epsilon}).$$

This immediately implies that

$$\|S_{J_*^c}\| \le O(\sqrt{n/\epsilon}) + O(\sqrt{n\log\epsilon^{-1}}) = O(\sqrt{n/\epsilon}).$$

In order to find μ and J_1 algorithmically, it suffices to apply a deterministic version of Theorem 3.7 due to Tropp [11, Theorem 3.1] to the matrices

$$(S - \frac{j}{\sqrt{n}}\mathbf{1}\mathbf{1}^T)_{J_0^c}, \quad j = 0, \pm 1, \pm 2, \ldots, \pm\lceil\sqrt{\log\epsilon^{-1}}\rceil.$$

By (5.1), Theorem 3.7, and the bound on $|\mathbb{E}\xi'|$, this procedure is guaranteed to succeed whenever (5.1) holds. We note that alternatively, one may proceed by first learning the mean $\mathbb{E}\xi'$ to within additive error $1/\sqrt{n}$ and then using the algorithm of [11, Theorem 3.1] on the matrix $(S - \widehat{\mathbb{E}S})_{J_0^c}$, where $\widehat{\mathbb{E}S}$ is our approximation for $\mathbb{E}S$.

Next, by symmetry, one may repeat this procedure to find a collection of $O(\epsilon n)$ rows K_* such that
$$\|S_{K_*^c}\| = O(\sqrt{n/\epsilon}).$$
By Lemma 4.1, it follows that zeroing out the $O(\epsilon n) \times O(\epsilon n)$ sub-matrix formed by the intersection of the rows in K_* and columns in J_* makes the norm of the resulting matrix $O(\sqrt{n/\epsilon})$. We now use Lemma 4.1 to combine the sub-matrices zeroed out for S, M_1, M_2, L and rescale ϵ to complete the proof.

For the running time of the algorithm, the algorithmic Grothendieck-Pietsch factorization in [11] can be performed in time $\tilde{O}(n^{7/2})$ and everything else can clearly be performed in time $\tilde{O}(n^2)$. $\square$

References

[1] Michael B. Cohen, *Nearly tight oblivious subspace embeddings by trace inequalities*, Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, ACM, New York, 2016, pp. 278–287.

[2] Uriel Feige and Eran Ofek, *Spectral techniques applied to sparse random graphs*, Random Structures Algorithms **27** (2005), 251–275.

[3] Wassily Hoeffding, *Probability inequalities for sums of bounded random variables*, J. Amer. Statist. Assoc. **58** (1963), 13–30.

[4] Can M. Le, Elizaveta Levina, and Roman Vershynin, *Concentration and regularization of random graphs*, Random Structures Algorithms **51** (2017), 538–561.

[5] Michel Ledoux and Michel Talagrand, *Probability in Banach spaces*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)], vol. 23, Springer-Verlag, Berlin, 1991, Isoperimetry and processes.

[6] Elizaveta Rebrova, *Constructive regularization of the random matrix norm*, J. Theoret. Probab. **33** (2020), 1768–1790.

[7] Elizaveta Rebrova and Roman Vershynin, *Norms of random matrices: local and global problems*, Adv. Math. **324** (2018), 40–83.

[8] Mark Rudelson and Roman Vershynin, *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.

[9] J. Schur, *Bemerkungen zur Theorie der beschränkten Bilinearformen mit unendlich vielen Veränderlichen*, J. Reine Angew. Math. **140** (1911), 1–28.

[10] Yoav Seginer, *The expected norm of random matrices*, Combinatorics, Probability and Computing **9** (2000), 149–166.

[11] Joel A. Tropp, *Column subset selection, matrix factorization, and eigenvalue optimization*, Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, PA, 2009, pp. 978–986.

[12] Y. Q. Yin, Z. D. Bai, and P. R. Krishnaiah, *On the limit of the largest eigenvalue of the large-dimensional sample covariance matrix*, Probab. Theory Related Fields **78** (1988), 509–521.

Simons Institute for the Theory of Computing, Berkeley, CA 94720, USA
Email address: `visheshj@stanford.edu`

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
Email address: `{asah,msawhney}@mit.edu`