ON THE SMOOTHED ANALYSIS OF THE SMALLEST SINGULAR VALUE WITH DISCRETE NOISE

VISHESH JAIN, ASHWIN SAH, AND MEHTAAB SAWHNEY

ABSTRACT. Let A be an $n \times n$ real matrix, and let M be an $n \times n$ random matrix whose entries are i.i.d sub-Gaussian random variables with mean 0 and variance 1. We make two contributions to the study of $s_n(A + M)$, the smallest singular value of $A + M$.

(1) We show that for all $\epsilon \geq 0$,
$$\mathbb{P}[s_n(A + M) \leq \epsilon] = O(\epsilon\sqrt{n}) + 2e^{-\Omega(n)},$$
provided only that A has $\Omega(n)$ singular values which are $O(\sqrt{n})$. This extends a well-known result of Rudelson and Vershynin, which requires all singular values of A to be $O(\sqrt{n})$.

(2) We show that any bound of the form
$$\sup_{\|A\| \leq n^{C_1}} \mathbb{P}[s_n(A + M) \leq n^{-C_3}] \leq n^{-C_2}$$
must have $C_3 = \Omega(C_1\sqrt{C_2})$. This complements a result of Tao and Vu, who proved such a bound with $C_3 = O(C_1 C_2 + C_1 + 1)$, and counters their speculation of possibly taking $C_3 = O(C_1 + C_2)$.

1. INTRODUCTION

Recall that for an $n \times n$ matrix A, its singular values $s_1(A) \geq \cdots \geq s_n(A)$ are defined to be the eigenvalues of $\sqrt{A^T A}$ arranged in non-decreasing order. Recall also that the largest and smallest singular values admit the following characterization:

$$s_1(A) = \sup_{x \in \mathbb{S}^{n-1}} \|Ax\|_2 = \|A\|;$$
$$s_n(A) = \inf_{x \in \mathbb{S}^{n-1}} \|Ax\|_2 = \|A^{-1}\|^{-1},$$

where $\mathbb{S}^{n-1}$ denotes the unit sphere in Euclidean space $\mathbb{R}^n$, $\|\cdot\|_2$ denotes the Euclidean norm in $\mathbb{R}^n$, and $\|\cdot\|$ denotes the spectral norm of the matrix.

The extreme singular values $s_1(A), s_n(A)$, and the condition number

$$\kappa(A) := s_1(A)/s_n(A) = \|A\| \cdot \|A\|^{-1}$$

are especially important in applications. In particular, consider the fundamental problem of solving the linear system $Ax = b$, where $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$ are given. Suppose that there is an error δb in inputting b, which leads to an error δx in the output. Then,

$$\frac{\|\delta x\|_2/\|x\|_2}{\|\delta b\|_2/\|b\|_2} = \left(\frac{\|A^{-1}(\delta b)\|_2}{\|\delta b\|_2}\right)\left(\frac{\|Ax\|_2}{\|x\|_2}\right) \leq \kappa(A),$$

which motivates defining the *loss of precision* [16] as $L(A) := \log \kappa(A)$.

1.1. **Smoothed analysis.** In their celebrated work on *smoothed analysis* [17], Spielman and Teng sought to understand why algorithms with poor worst-case performance are successful in practice. In the context of solving the linear system $Ax = b$, their program amounts to the following: even if the desired input matrix A is ill-conditioned (i.e, has a large condition number, and hence, large loss of precision), it is likely that the computer will actually work with some small perturbation

$A + M$, where M represents the effect of random noise. Therefore, if we can show that for any fixed $A \in \mathbb{R}^{n \times n}$, the random matrix $A + M$ is well-conditioned (i.e., has small condition number) with high probability, this would help explain why the large condition number of A does not lead to correspondingly bad performance in practice.

Since the operator norm of a random matrix M with independent entries (with sufficiently light tails) can be controlled using standard concentration techniques (cf. [2,22]) (this immediately implies control on the operator norm of $A + M$ using the triangle inequality), it follows that they key challenge in controlling the condition number of $A + M$ is in understanding the behavior of the smallest singular value of $A + M$. In this direction, Sankar, Spielman, and Teng [15] considered the case when the entries of M are i.i.d Gaussian random variables with mean 0 and variance 1, and showed that for any $A \in \mathbb{R}^{n \times n}$ and for any $\epsilon \geq 0$,

$$\mathbb{P}[s_n(A + M) \leq \epsilon] \leq C\epsilon\sqrt{n}, \tag{1.1}$$

for $C \sim 2.35$ (recently, the constant C was improved by Banks et al. to the optimal value of 1 in [3]). This result is almost best possible, in the sense that when $A = 0$, it was shown by Edelman [5] that for sufficiently small $\epsilon \geq 0$,

$$\mathbb{P}[s_n(M) \leq \epsilon] \geq c\epsilon\sqrt{n},$$

where $c > 0$ is an absolute constant (in fact, c can be taken arbitrarily close to 1 by further restricting the range of ϵ).

1.2. **Smallest singular value of random matrices.** The proof of Sankar, Spielman, and Teng heavily exploited the orthogonal invariance of the Gaussian distribution, and extending their result to other distributions, especially discrete distributions, is very challenging, and has attracted much attention in the past fifteen years. Notably, in a seminal work, Rudelson and Vershynin [12] showed that for a random matrix M whose entries are i.i.d copies of a sub-Gaussian random variable with variance 1 and mean 0 (recall that a random variable X is said to be sub-Gaussian if its sub-Gaussian norm, defined by $\|X\|_{\psi_2} := \inf\{s > 0 : \mathbb{E}[\exp(X^2/s^2)] \leq 2\}$, is finite),

$$\mathbb{P}[s_n(M) \leq \epsilon] \leq C\epsilon\sqrt{n} + 2\exp(-cn), \tag{1.2}$$

where $C, c > 0$ depend only on the sub-Gaussian norm of the distribution. We note that this result is almost best-possible; the first summand has the optimal dependence on ϵ and n as explained before, whereas the second summand is necessary to account for the probability that the random matrix is singular (for instance, if each entry of M is independently ± 1 with probability $1/2$ each, then any two given rows are equal with probability 2^{-n}).

Since the appearance of [12], considerable work has been devoted to relaxing the assumptions in the work of Rudelson and Vershynin, and establishing a bound of the form (1.2) in greater generality. For instance, it was very recently shown by Livshyts, Tikhomirov, and Vershynin [8] (see also the references there for previous work) that (1.2) holds under only the assumption that the entries of M are independent, have mean 0, have uniformly bounded variance, and are uniformly anti-concentrated (in fact, their result is a bit more general, as we will discuss later). In the setting of smoothed analysis, Tikhomirov [20] showed that for any $A \in \mathbb{R}^{n \times n}$ and for all $\epsilon > 0$,

$$\mathbb{P}[s_n(A + M) \leq \epsilon] \leq C\epsilon\sqrt{n},$$

if the rows of M are independent *continuous* random vectors satisfying certain technical assumptions (in particular, his result allows the rows of M to be independent centered log-concave isotropic random vectors). We mention that when the entries of M are independent continuous random variables with uniformly bounded densities, a bound of the form (1.1), with the optimal $\sqrt{n}$ replaced by the sub-optimal n, is rather easy to prove (cf. [20]); for *symmetric* matrices M whose distributions on and above the diagonal are independent continuous random variables with uniformly bounded densities and *symmetric* matrices A, a bound of the form (1.1) was obtained by Farrell and Vershynin

[6], except that $\sqrt{n}$ is replaced by the sub-optimal n^2 (in the special case when M is drawn from the appropriately normalised Gaussian Orthogonal Ensemble, the optimal $\sqrt{n}$ dependence was obtained by Bourgain [4]).

1.3. **Smoothed analysis with discrete noise.** Since the noise/randomness encountered in practice (e.g., in digital devices) is typically discrete, it is natural to try to understand analogues of (1.1) when the distribution of the entries of M is allowed to have atoms. This was first considered by Tao and Vu [18], who showed that if the entries of M are i.i.d copies of a (possibly complex) random variable with mean 0 and variance 1, then for any constants $C_1, C_2 > 0$, there exists a $C_3 > 0$ such that

$$\sup_{\|A\| \leq n^{C_1}} \mathbb{P}[s_n(A + M) \leq n^{-C_3}] \leq n^{-C_2}. \tag{1.3}$$

Recently, it was shown by Jain [7] that a bound of this nature continues to hold under the more general condition $\|A\| \leq \exp(n^c)$, and for all target probabilities larger than $\exp(-n^{c'})$, where $c, c' > 0$ are small absolute constants. The best-known dependence of C_3 on C_1, C_2 was obtained in another work by Tao and Vu [19], who showed that

$$\sup_{\|A\| \leq n^{C_1}} \mathbb{P}[s_n(A + M) \leq n^{-2(C_2+2)C_1+1/2}] \leq O(n^{-C_2+o(1)} + \mathbb{P}[\|M\| \geq n^{C_1}]). \tag{1.4}$$

Perhaps surprisingly, it turns out that the dependence of the smallest singular value on C_1 is not an artifact of the proof, and is necessary in general, as the following unpublished example due to Rudelson shows (a similar example appears in independent work of Tao and Vu [19]).

Example 1.1. Let $k \in [n]$, and let $A = \mathrm{diag}(L, L, \ldots, L, 0, 0, \ldots, 0)$, where the first $n - k$ entries are L, and the last k entries are 0. Let the entries of M be independent lazy Rademacher random variables (i.e., each entry independently takes on the value 0 with probability $1/2$, -1 with probability $1/4$, and 1 with probability $1/4$). Let $(v_1, \ldots, v_n) \in \mathbb{S}^{n-1}$ be orthogonal to the first $n - 1$ rows of $A + M$. It is easily seen that $\sqrt{v_1^2 + \cdots + v_{n-k}^2} = O(\sqrt{n}/L)$. Then, since the last k coordinates of the last row are 0 with probability 2^{-k}, it follows that with probability at least 2^{-k}, the smallest singular value of $A + M$ is at most $O(\sqrt{n}/L)$.

1.4. **Our results.** Note that Example 1.1 rules out the possibility of a bound of the form (1.2) for general $A + M$, even when the entries of M are i.i.d sub-Gaussian random variables. On the other hand, such a bound is known to hold (restricting ourselves to the case when M has i.i.d sub-Gaussian entries with mean 0 and variance 1) if $\|A\| = O(\sqrt{n})$ [12], and more generally, if $\sum_{i=1}^{n} s_i(A)^2 = O(n^2)$ [8].

Our first main result shows that when the entries of M are i.i.d sub-Gaussian random variables with mean 0 and variance 1, (1.2) continues to hold for a much wider class of $A + M$. In order to state the result, we need the following definition.

Definition 1.2. Let $m \leq n \in \mathbb{N}$ and $K \geq 0$. An $m \times n$ matrix A has K-rank r if it has exactly r singular values of size at least $K\sqrt{n}$.

Theorem 1.3. *Let $K > 0$ and $\eta \in (0, 1)$. Let ξ be a sub-Gaussian random variable with mean 0 and variance 1, and let M be an $n \times n$ random matrix, each of whose entries is an independent copy of ξ. Then, for any $A \in \mathbb{R}^{n \times n}$ with K-rank at most $(1 - \eta)n$, we have*

$$\mathbb{P}[s_n(A + M) \leq \epsilon] \leq C\epsilon\sqrt{n} + 2e^{-cn}, \tag{1.5}$$

where $C, c > 0$ depend only on K, η, and the sub-Gaussian norm of ξ.
Remark.

3

(1) In the setting where M has i.i.d sub-Gaussian entries with mean 0 and variance 1, the condition that the K-rank of A is at most $(1-\eta)n$ is a relaxation of the condition $\sum_{i=1}^{n} s_i(A)^2 = O(n^2)$ appearing in [8] (although, note that the noise M in [8] may be inhomogeneous, and is allowed to have much heavier tails). Notably, it encompasses and generalizes the class of low-rank matrices, which are especially important in applications (see, e.g., [1, 10]).

(2) The dependence of c on η is necessary in general. Indeed, Example 1.1 provides an example of A and M such that A has 0-rank $n-k$, and such that on an event of probability at least 2^{-k}, $s_n(A+M)$ depends on $\|A\|$.

(3) The only place where sub-Gaussianity of ξ is used in our proof is to guarantee that $\|M\| = O(\sqrt{n})$ except with exponentially small probability. It is well-known [2] that $\|M\| = O(\sqrt{n})$ with high probability, as long as the entries of M have uniformly bounded fourth moment. Consequently, Theorem 1.3 continues to hold under the more general assumption that ξ is a random variable with mean 0, variance 1, and finite fourth moment, provided we add the term $(1 - \mathbb{P}[\|M\| = O(\sqrt{n})])$ to the right-hand side of (1.5).

(4) In [21], it was shown by Tikhomirov that if M' is an $N \times n$ matrix ($N \geq (1+\delta)n$) with i.i.d uniformly anti-concentrated entries, then for any $N \times n$ matrix A', $\mathbb{P}[s_n(A' + M') \leq u\sqrt{N}] \leq 2\exp(-vN)$, where u, v depend only on $\delta > 0$ and the uniform bound on the anti-concentration. Compared to [21], the main innovation in our work (Lemma 2.14) is the consideration of the arithmetic structure of normal vectors to random hyperplanes without any *a priori* control on the operator norm (this is unnecessary in the rectangular case).

Next, we study the influence of A on the lower tail of $s_n(A+M)$. As noted above, in the case when $\|A\| \leq n^{C_1}$ and the target probability is n^{-C_2}, the best-known lower bound on the smallest singular value due to Tao and Vu (1.4) is of the form $n^{-O(C_1 C_2 + C_1 + 1)}$. It is natural to ask whether this dependence may be improved, in particular, if the term $C_1 C_2$ may be replaced by $C_1 + C_2$ (this seems to have been speculated by Tao and Vu in [19]). Our second main result shows that a dependence of the form $O(C_1 + C_2)$ is not possible, and in fact, that the term $C_1 C_2$ cannot be replaced by anything asymptotically smaller than $C_1 \sqrt{C_2}$.

Theorem 1.4. *Let R_n be an $n \times n$ random matrix, each of whose entries is an independent copy of a lazy Rademacher random variable (i.e., 0 with probability $1/2$, and ± 1 with probability $1/4$ each). There exist positive constants $K, C > 0$ for which the following holds. Let $L \geq 2K\sqrt{n}$, and let $A = \mathrm{diag}(L, \ldots, L, 0)$. Then, for each positive integer t and for all n sufficiently large (depending only on t), we have*

$$\mathbb{P}\left[s_n(A+R_n) \leq C\left(\frac{K\sqrt{n}}{L}\right)^t\right] \geq (4C)^{-t} K^{-\frac{(t-1)(t-2)}{2}} (\log 2t)^{1-t} n^{-\frac{t(t-1)}{4}}.$$

Remark. The dependence of the probability on n should be sharp. However, establishing this would essentially require a joint local central limit theorem involving certain low-degree polynomials. Moreover, such a result would likely allow extension to the case when the entries of R_n are Rademacher random variables (as we will see, our proof is considerably simplified by using the 2-divisibility of lazy Rademacher random variables).

2. Proof of Theorem 1.3

The proof of Theorem 1.3 follows the (by now) standard geometric framework pioneered by Rudelson and Vershynin [12]. The key deviation in our argument is in showing that a unit normal vector to the random hyperplane spanned by the (say) first $n-1$ rows of the matrix is arithmetically

very unstructured (in the sense of having exponentially large Least Common Denominator (LCD)). The union bound argument in [12] or its refinement in [8] is inadequate here since the matrix A may be arbitrarily large. Our innovation (Lemma 2.14) is to execute a union bound argument by covering level sets of the LCD by certain oblique convex sets adapted to the geometry of A, which exploits the fact that it suffices to cover only those vectors v which are "approximately orthogonal" to the large singular vectors of A.

Throughout this section, we will fix K and η as in the statement of Theorem 1.3, and let B and N denote the $(n-1) \times n$ matrices consisting of the first $n-1$ rows of A and M respectively. Note that the B has at least $\eta n/2$ singular values smaller than $K\sqrt{n}$, as is readily seen by applying the Cauchy interlacing theorem to the matrix AA^T and its minor BB^T. Throughout, we will let $V \subseteq \mathbb{R}^{n-1}$ denote the span of the bottom $\eta n/2$ left-singular vectors of B (note that the singular values corresponding to these singular vectors are all at most $K\sqrt{n}$). Moreover, P_V will denote the orthogonal projection operator from $\mathbb{R}^{n-1}$ onto V.

We will use the decomposition of the unit sphere into compressible and incompressible vectors, formalized by Rudelson and Vershynin [12].

Definition 2.1 (Compressible and Incompressible vectors, [12, Definition 3.2]). Fix $\delta, \rho \in (0,1)$. A vector $x \in \mathbb{R}^N$ is said to be *sparse* if $|\operatorname{supp}(x)| \leq \delta n$. A vector $x \in \mathbb{S}^{N-1}$ *compressible* if it is within Euclidean distance ρ of a sparse vector, and *incompressible* otherwise. We denote the set of compressible and incompressible vectors by $\operatorname{Comp}_N(\delta, \rho)$ and $\operatorname{Incomp}_N(\delta, \rho)$ respectively, dropping the subscript N when the underlying dimension is clear from context.

We will also need the notion of the Lévy concentration function.

Definition 2.2 (Lévy concentration function). For a random variable X and a real number $r \geq 0$, we define the Lévy concentration function of X at radius r by

$$\mathcal{L}(X, r) = \sup_{y \in \mathbb{R}} \mathbb{P}[|X - y| \leq r].$$

First, we show that any fixed vector $v \in \mathbb{S}^{n-1}$ has exponentially small probability of being in the kernel of the first $n-1$ rows of $(A + M)$. In fact, we will show something stronger: for every fixed vector $v \in \mathbb{S}^{n-1}$, its image under $P_V(B+N)$ has norm $\Omega(\sqrt{n})$ except with exponentially small probability; this strengthening will be crucial for the union bound argument in Lemma 2.4.

The next lemma follows immediately from [21, Corollary 6], which is a direct consequence of the main result in [14]. For sub-Gaussian random variables, it follows easily from the Hanson-Wright inequality (straightforwardly modifying the proofs of [13, Corollary 2.4] and [13, Corollary 3.1]); we include the deduction for the reader's convenience.

Lemma 2.3 (Invertibility on a single vector). *There exists $c_{2.3} > 0$ depending only on η and the sub-Gaussian norm of ξ such that for any $v \in \mathbb{S}^{n-1}$,*

$$\mathbb{P}[\|P_V(B+N)v\|_2 \leq c_{2.3}\sqrt{n}] \leq 2e^{-c_{2.3}n}.$$

Proof. Let N and N' be independent copies of N. Then

$$\mathbb{P}[\|P_V(B+N)v\|_2 \leq \sqrt{\eta n}/4]^2 = \mathbb{P}[\|P_V(B+N)v\|_2 \leq \sqrt{\eta n}/4 \cap \|P_V(B+N')v\|_2 \leq \sqrt{\eta n}/4]$$
$$\leq \mathbb{P}[\|P_V(N - N')v\|_2 \leq \sqrt{\eta n}/2].$$

Note that the coordinates of $(N - N')$ are i.i.d sub-Gaussian random variables with variance 2, and sub-Gaussian norm depending only on that of ξ. In particular,

$$\mathbb{E}[\|P_V(N - N')v\|_2^2] = 2\operatorname{Tr}(P_V^T P_V) \geq \eta n.$$

5

The result now follows from the Hanson-Wright inequality applied to the matrix P_V and vector $(N - N')v$ (see [13, Theorem 2.1]), which shows that

$$\mathbb{P}[|\|P_V(N - N')v\|_2 - 2\operatorname{Tr}(P_V^T P_V)| > t] \leq 2\exp(c_\xi t^2),$$

where c_ξ depends only on the sub-Gaussian norm of ξ. $\qquad\square$

Given this, we can quickly derive invertibility on compressible vectors.

Lemma 2.4 (Invertibility on compressible vectors). *There exist $\delta_{2.4}, \rho_{2.4}, c_{2.4} > 0$ depending only on η, K, and the sub-Gaussian norm of ξ for which*

$$\mathbb{P}[\exists v \in \operatorname{Comp}_{\delta_{2.4}, \rho_{2.4}} : \|(B + N)v\|_2 \leq c_{2.4}\sqrt{n}] \leq 2e^{-c_{2.4}n}.$$

Proof. The key point is that, by definition of V, $\|P_V B\| \leq K\sqrt{n}$. Since ξ has bounded sub-Gaussian norm, it follows (cf. [12, Lemma 2.4]) that there exists K', depending only on K and the sub-Gaussian norm of ξ, for which the event $\mathcal{E}_{K'} = \{\|P_V(B + N)\| \leq K'\sqrt{n}\}$ occurs with probability at most $2e^{-cn}$, for some $c > 0$ depending only on the sub-Gaussian norm of ξ.

Let $\delta, \rho, c' > 0$, and consider an ϵ-net $\mathcal{N}$ of $\operatorname{Comp}_{\delta, \rho}$. Then, we have

$$\mathbb{P}[\exists v \in \operatorname{Comp}_{\delta, \rho} : \|(B + N)v\|_2 \leq c'\sqrt{n}]$$
$$\leq \mathbb{P}[\mathcal{E}_{K'}] + \mathbb{P}[\exists v \in \operatorname{Comp}_{\delta, \rho} : \|P_V(B + N)v\|_2 \leq c'\sqrt{n} \wedge \mathcal{E}_{K'}^c]$$
$$\leq 2e^{-cn} + \mathbb{P}[\exists v \in \mathcal{N} : \|P_V(B + N)v\|_2 \leq (c' + K'\epsilon)\sqrt{n}].$$

Therefore, by using Lemma 2.3, we can conclude, provided we first choose c', ϵ small enough so that $c' + K'\epsilon < c_{2.3}$ (this is clearly possible), and then choose δ, ρ small enough so that there is an ϵ-net $\mathcal{N}$ of $\operatorname{Comp}_{\delta, \rho}$ of size at most $\exp(c_{2.3}n/2)$ (this is possible by a standard volumetric bound on the size of ϵ-nets of $\operatorname{Comp}_{\delta, \rho}$, cf. [11, Lemma 4.3]). We omit the standard details. $\qquad\square$

For invertibility on incompressible vectors, we will use the following reduction due to Rudelson and Vershynin.

Lemma 2.5 (Invertibility via distance, [12, Lemma 3.5]). *Let A be a random $n \times n$ matrix. Let X_i denote the column vectors of A, and let $H_k = \operatorname{span}(X_{-k})$ denote the span of all columns except for the k^{th} column. Then for every $\delta, \rho \in (0, 1)$ and $\epsilon \geq 0$ we have that*

$$\mathbb{P}\left[\inf_{x \in \operatorname{Incomp}(\delta, \rho)} \|Ax\|_2 \leq \epsilon\rho n^{-1/2}\right] \leq \frac{1}{\delta n}\sum_{k=1}^n \mathbb{P}[\operatorname{dist}(X_k, H_k) \leq \epsilon].$$

Given this, our goal is to find a uniform (in $k \in [n]$) upper bound on $\mathbb{P}[\operatorname{dist}(X_k, H_k) \leq \epsilon]$.

We will need the crucial notion of the Least Common Denominator due to Rudelson and Vershynin.

Definition 2.6 (Least common denominator, cf. [11]). For a vector $v \in \mathbb{R}^N$, $\gamma \in (0, 1)$, and $\alpha > 0$, we define

$$\operatorname{LCD}_{\alpha, \gamma}(v) = \inf\{\theta > 0 : \operatorname{dist}(\theta v, \mathbb{Z}^N) < \min(\gamma\|\theta v\|, \alpha)\}.$$

We collect some useful properties of the LCD.

Lemma 2.7 (Lower bound on LCD, cf. [11, Lemma 6.1]). *Fix $\delta, \rho > 0$. There exist $\gamma, \lambda > 0$ (depending only on δ, ρ) such that for any $v \in \operatorname{Incomp}_N(\delta, \rho)$ and for all $\alpha > 0$, we have $\operatorname{LCD}_{\alpha, \gamma}(v) \geq \lambda\sqrt{n}$.*

Lemma 2.8 (Anti-concentration via LCD). *For any $\gamma \in (0,1)$, there exist $c_{2.8}(\gamma), C_{2.8}(\gamma) > 0$ depending only on γ and the sub-Gaussian norm of ξ for which the following holds. Let $v \in \mathbb{S}^{N-1}$. Then, for every $\alpha > 0$ and $\epsilon \geq 0$,*

$$\mathcal{L}\left(\sum_{i=1}^{n} v_i \xi_i, \epsilon\right) \leq C_{2.8}(\gamma)\epsilon + \frac{C_{2.8}(\gamma)}{\mathrm{LCD}_{\alpha,\gamma}(v)} + C_{2.8}(\gamma)e^{-c_{2.8}(\gamma)\alpha^2},$$

where $\xi_1, \ldots, \xi_n$ denote i.i.d copies of ξ.

Given Lemma 2.8, a good upper bound on $\mathbb{P}[\mathrm{dist}(X_k, H_k) \leq \epsilon]$ may be obtained by showing that it is exponentially unlikely that any unit normal to H_k has subexponential LCD. The proof of Rudelson and Vershynin in the centered sub-Gaussian case accomplishes this by decomposing the set of incompressible vectors into level sets of the LCD (Definition 2.9), and then finding, for each level set, a net (at appropriate scale) of small enough size to survive a union bound argument. We note that this step requires considerable care, and there is not much room available in the entropy-energy trade-off.

Definition 2.9 (Level sets of LCD). Let $D, \gamma, \mu > 0$. We define

$$S_D = \{x \in \mathbb{S}^{n-1} : D \leq \mathrm{LCD}_{\mu\sqrt{n},\gamma}(x) \leq 2D\}.$$

In our setting, such a net argument for S_D has no hope of working, since the operator norm of $A + M$ can be arbitrarily large. The key idea to overcome this is to exploit the fact that for a sufficiently large constant K', the following holds except with exponentially small probability: $\|Bv\|_2 > 2K'\sqrt{n}$ implies that $\|(B + N)v\|_2 > K'\sqrt{n}$. This motivates the following definition.

Definition 2.10 (Restricted level sets of the LCD). For $K' > 0$, let

$$G_{K'} = \{\|x\|_2 \leq 1 \cap \|Bx\|_2 \leq 2K'\sqrt{n}\},$$

and let

$$S'_D(K') = S_D \cap G_{K'}.$$

In order to find an efficient net for $S'_D(K')$, we will use the following two easy geometric observations repeatedly. Note that the remainder of the proofs, parallelepipeds will always be "right-parallelepipeds" or more colloquially "boxes".

Lemma 2.11. *There exists an absolute constant $C_{2.11} \geq 1$ such that any ellipsoid with semiaxes ℓ_i can be covered by at most $C_{2.11}^n$ parallelepipeds with widths $\ell_i/\sqrt{n}$ in the direction of the semiaxes.*

Proof. This follows from a standard volumetric argument. $\square$

Lemma 2.12. *The sets $S_1 = \{\|x\|_2 \leq 1 \cap \|Jx\|_2 \leq 1\}$ and $S_2 = \{\|x\|_2^2 + \|Jx\|_2^2 \leq 1\}$ satisfy $S_2 \subseteq S_1 \subseteq \sqrt{2}S_2$.*

Proof. This is immediate. $\square$

We will also need the following elementary lattice counting fact.

Lemma 2.13. *The number of integer lattice points in a right-parallelepiped with dimensions $\ell_i \geq 1$ is at most $C_{2.13}^n \prod_{i=1}^{n} \ell_i$, where $C_{2.13} \geq 1$ is an absolute constant.*

Remark. Note the parallelepiped does not need to be axis-aligned or centered.

Proof. By expanding the dimensions to $\lceil \ell_i \rceil \leq 2\ell_i$, and by decomposing the region into unit cubes whose axes are aligned with those of the parallelepiped, it suffices to show that there is an absolute constant $C \geq 1$ such that any rotated unit cube has at most C^n integer lattice points in it. Since any rotated unit cube is contained in a ball of radius $\sqrt{n}$, it suffices to show that any ball (not necessarily centered) of radius $\sqrt{n}$ contains at most C^n integer lattice points, for an absolute constant $C \geq 1$, which follows from a standard well-known volumetric argument. $\square$

We are now ready to state and prove our key new ingredient.

Lemma 2.14 (Nets of level sets of LCD). *Fix $\lambda, K' > 0$. There is $C_{2.14} = C_{2.14}(\lambda, K, K')$ for which the following holds. For all sufficiently small $\mu > 0$ (depending on λ) and for all $D \geq \lambda\sqrt{n}$, there exists a net $\mathcal{N}$ of cardinality at most $\mu^{-(1-\eta/2)n}D^2(C_{2.14}D/\sqrt{n})^n$ with the following property: on the event $\|N\| \leq K'\sqrt{n}$, for any $v \in S'_D(K')$, there is $w \in \mathcal{N}$ satisfying*

$$\|(B+N)(v-w)\|_2 \leq \frac{\mu n}{D}.$$

Proof. For lightness of notation, we will denote $S'_D(K')$ by S'_D. We denote the singular values of B, in decreasing order, by $\sigma_1(B), \ldots, \sigma_n(B)$, and the corresponding unit right-singular vectors by $b_1, \ldots, b_n$. In particular, $b_1, \ldots, b_n$ form an orthonormal basis of $\mathbb{R}^n$. Note that since B is an $(n-1) \times n$ matrix, we must have $\sigma_n(B) = 0$, and so b_n is in the kernel of B. The proof will make use of a few different geometric objects; we collect their definitions here for ease of reference:

- $\mathcal{C}$ denotes the n-dimensional cube of width μ/D centered at 0, whose axes are aligned with $b_1, \ldots, b_n$.
- $\mathcal{G}$ denotes the region given by $\{x \in \mathbb{B}_2^n : \|Bx\|_2 \leq 2K'\sqrt{n}\}$.
- $\mathcal{Q}$ denotes the ellipsoid given by $\{4K'^2 n\|x\|_2^2 + \|Bx\|_2^2 \leq 8K'^2 n\}$.
- $\mathcal{C}'$ denotes the right-parallelepiped centered at 0 whose axes are aligned with $b_1, \ldots, b_n$, and whose width along b_i is

$$\sqrt{\frac{8K'^2}{\sigma_i(B)^2 + 4K'^2 n}};$$

  note that $\mathcal{C}' \subseteq \mathcal{Q}$.
- $\mathcal{C}''$ denotes the right-parallelepiped centered at 0 whose axes are aligned with $b_1, \ldots, b_n$, and whose width along b_i is

$$\sqrt{\frac{8K'^2}{\sigma_i(B)^2 + 4K'^2 n}} + \frac{1}{D};$$

  note that $\mathcal{C}' + \mu^{-1}\mathcal{C} \subseteq \mathcal{C}''$.
- $\mathcal{C}'''$ denotes the right-parallelepiped centered at 0 whose axes are aligned with $b_1, \ldots, b_n$, and whose width along b_i is

$$\frac{1}{\min(D, \max(\sigma_i(B), \sqrt{n}))}$$

Step 1: From the definition of LCD, it follows that S_D admits a $(2\mu\sqrt{n}/D)$-net in Euclidean norm, formed by the points

$$\mathcal{P} = \left\{\frac{p}{\|p\|_2} : p \in (\mathbb{Z}^n \setminus \{0\}) \cap \mathbb{B}_2^n(0, 3D)\right\} = \left\{\frac{p}{\|p\|_2} : p \in (\mathbb{Z}^n \setminus \{0\}) \cap (\mathbb{B}_2^n(0, 3D) \setminus \mathbb{B}_2^n(0, 3D/2))\right\}.$$

For a concrete reference for the first equality, see [11, Lemma 7.2]. For the second equality, simply note that for any $p \in (\mathbb{Z}^n \setminus \{0\}) \cap \mathbb{B}_2^n(0, 3D/2)$, there exists some $\ell_p \in \mathbb{Z} \setminus \{0\}$ such that $\ell_p \cdot p \in (\mathbb{Z}^n \setminus \{0\}) \cap (\mathbb{B}_2^n(0, 3D) \setminus \mathbb{B}_2^n(0, 3D/2))$. Since $\ell_p \cdot p/\|\ell_p \cdot p\|_2 = p/\|p\|_2$, both p and $\ell_p \cdot p$ map to the same point in $\mathcal{P}$, so that we may ignore the contribution of p without any loss.

Hence, we see that $S_D \subseteq \mathcal{P} + \mathbb{B}_2^n(0, 2\mu\sqrt{n}/D)$, so that

$$S'_D \subseteq (\mathcal{P} + \mathbb{B}_2^n(0, 2\mu\sqrt{n}/D)) \cap \mathcal{G}.$$

In the remainder of the proof, we will cover the region on the right hand side by at most

$$\mu^{-(1-\eta/2)n}D^2(CD/\sqrt{n})^n$$

8

translates of

$$\frac{\mu\sqrt{n}}{4K'D}\mathcal{Q};$$

this clearly suffices since on the event $\|N\| \le K'\sqrt{n}$, for any $v, w \in \mu\sqrt{n}\mathcal{Q}/4K'D$, we have $\|(B + N)(v - w)\|_2 \le \mu n/D$.

Step 2: Since for any $p \in (\mathbb{Z}^n \setminus \{0\}) \cap \mathbb{B}_2^n(0, 3D)$, $\|p\|_2^2 \in \mathbb{Z} \cap [1, 9D^2]$, it follows that by paying an overall multiplicative factor of $9D^2$ in the size of the final net, it suffices to fix $T \in [3D/2, 3D]$ and bound (uniformly in T) the number of translates of $\mu\sqrt{n}\mathcal{Q}/4K'D$ needed to cover the region

$$(\mathcal{P}_T + \mathbb{B}_2^n(0, 2\mu\sqrt{n}/D)) \cap \mathcal{G},$$

where

$$\mathcal{P}_T = \left(\frac{1}{T}\mathbb{Z}^n\right) \cap \mathbb{B}_2^n.$$

Moreover, since $\mathbb{B}_2^n(0, 2\mu\sqrt{n}/D)$ can be covered by $(2C_{2.11})^n$ translates of $\mathcal{C}$ (Lemma 2.11), it suffices after paying a multiplicative factor of $(2C_{2.11})^n$ to bound (uniformly in T and $y \in \mathbb{R}^n$) the number of translates of $\mu\sqrt{n}\mathcal{Q}/4K'D$ needed to cover the region

$$(y + \mathcal{P}_T + \mathcal{C}) \cap \mathcal{G}.$$

Moreover, by Lemma 2.12, it suffices instead to consider the larger region

$$(y + \mathcal{P}_T + \mathcal{C}) \cap \mathcal{Q}.$$

Note that covering $\mathcal{Q}$ by translates of $\mu\sqrt{n}\mathcal{Q}/4K'D$ requires $\mu^{-n}(CD/\sqrt{n})^n$ translates, which is bigger by a factor of $\mu^{\eta n/2}$ compared to our desired conclusion; therefore, we must carefully exploit the first term in the intersection.

Step 3: By noting that the matrix $(4K'^2 nI + B^T B)/(8K'^2 n)$ has unit right-eigenvectors $b_1, \ldots, b_n$ with corresponding eigenvalues $(4K'^2 n + \sigma_i(B)^2)/8K'^2 n$, it follows from Lemma 2.11 that $\mathcal{Q}$ can be covered by $C_{2.11}^n$ translates of $\mathcal{C}'$. Therefore, up to an overall multiplicative factor of $C_{2.11}^n$, it suffices to bound (uniformly in T, and $y_1, y_2 \in \mathbb{R}^n$) the number of translates of $\mu\sqrt{n}\mathcal{Q}/4K'D$ needed to cover the region

$$(y_1 + \mathcal{P}_T + \mathcal{C}) \cap (y_2 + \mathcal{C}').$$

Moreover, since $\mathcal{C}' \subseteq \mathcal{Q}$, it suffices to cover by translates of $\mu\sqrt{n}\mathcal{C}'/4K'D$. Below, we will need the following observation. Let

$$k^* = \max\{i \in [n] : \sigma_i(B) \ge KD/\lambda\}.$$

Then, since $D \ge \lambda\sqrt{n}$ by assumption, and recalling the observation that B has K-rank at most $(1 - \eta/2)n$, it follows that $k^* \le (1 - \eta/2)n$.

Step 4: We begin by bounding (from above) the number of points $w \in y_1 + \mathcal{P}_T$ for which

$$(w + \mathcal{C}) \cap (y_2 + \mathcal{C}') \ne \emptyset.$$

Note that any such w is of the form $y_1 + (z/T)$, where $z \in \mathbb{Z}^n$, and for the intersection above to be nonempty, we must have

$$(z/T) \in (y_2 - y_1) + \mathcal{C}' - \mathcal{C} \subseteq (y_2 - y_1) + \mathcal{C}''.$$

We claim that $\mathcal{C}'' \subseteq C \cdot \mathcal{C}'''$, where C is a constant depending only on K, K', λ. Indeed, using the definition of $\mathcal{C}'', k^*$, and the assumption $D \ge \lambda\sqrt{n}$, we see that the width of $\mathcal{C}''$ in the directions b_i for $i \le k^*$ is at most C''/D, whereas its width in the directions b_i for $i > k^*$ is at most $C''/\max(\sigma_i(B), \sqrt{n})$, where C'' is a constant depending on K, K', λ.

Hence, it suffices to bound the number of $z \in \mathbb{Z}^n$ such that $z \in T \cdot (y_2 - y_1) + T \cdot C \cdot C'''$. Moreover, since $T \leq 3D$, it suffices to bound (uniformly in $y \in \mathbb{R}^n$) the number of $z \in \mathbb{Z}^n$ such that $z \in y + 3CD \cdot C'''$. Since for all sufficiently large C (depending only on K, K', λ),

$$\frac{3CD}{\min(D, \max(\sqrt{n}, \sigma_i(B)))} \geq 1,$$

it follows from Lemma 2.13 that the number of such $z \in \mathbb{Z}^n$ is at most

$$N_1 := \frac{(CD)^n}{\prod_{i=1}^n (\min(D, \max(\sqrt{n}, \sigma_i(B))))},$$

where C is a constant depending only on K, K', λ.

Step 5: Let us now fix $w_0 \in y_1 + \mathcal{P}_{\mathcal{T}}$ and bound (uniformly in $y_1, y_2 \in \mathbb{R}^n$) the number of translates of $\mu\sqrt{n}\mathcal{C}'/4K'D$ needed to cover the region

$$(w_0 + \mathcal{C}) \cap (y_2 + \mathcal{C}').$$

For this, note that for any $y_2 \in \mathbb{R}^n$, $(w_0 + \mathcal{C}) \cap (y_2 + \mathcal{C}')$ is the intersection of two right-parallelepipeds along the same axes $b_1, \ldots, b_n$. In particular, in each direction b_i, the width of $(w_0 + \mathcal{C}) \cap (y_2 + \mathcal{C}')$ is bounded above by the minimum of the widths of $\mathcal{C}$ and $\mathcal{C}'$ along b_i. We will use the bound for the width coming from $\mathcal{C}'$ for the directions b_i, $i \leq k^*$, and the bound for the width coming from $\mathcal{C}$ for the remaining directions.

Thus, the width of $\mu\sqrt{n}\mathcal{C}'/4K'D$ is smaller than the width of $(w_0 + \mathcal{C}) \cap (y_2 + \mathcal{C}')$ by a factor of at most $\mu\sqrt{n}/4K'D$ in the directions b_i for $i \leq k^*$. Moreover, for $i > k^*$, the width of $(w_0 + \mathcal{C}) \cap (y_2 + \mathcal{C}')$ in direction b_i is at most μ/D, whereas the width of $\mu\sqrt{n}\mathcal{C}'/4K'D$ is at least $\mu\sqrt{n}/C'D \max(\sigma_i(B), \sqrt{n})$, where C' is a constant depending only on K, K', λ; hence, in these directions, the width of $\mu\sqrt{n}\mathcal{C}'/4K'D$ is smaller by a factor of at most $\sqrt{n}/C' \max(\sigma_i(B), \sqrt{n})$.

Therefore, we see that the number of translates of $\mu\sqrt{n}\mathcal{C}'/4K'D$ needed to cover $(w_0 + \mathcal{C}) \cap (y_2 + \mathcal{C}')$ is at most

$$N_2 := C^n \prod_{i=1}^{k^*} \left(\frac{D}{\mu\sqrt{n}}\right) \prod_{i=k^*+1}^n \left(\frac{\max(\sigma_i(B), \sqrt{n})}{\sqrt{n}}\right),$$

where C is a constant depending only on K, K', λ.

Step 6: Noting that

$$\prod_{i=1}^n \min(D, \max(\sqrt{n}, \sigma_i(B))) \geq (C')^{-n} D^{k^*} \prod_{k^*+1}^n \max(\sigma_i(B), \sqrt{n}),$$

where C' is a constant depending only on K, K', λ, we see that

$$N_1 \cdot N_2 \leq \mu^{-k^*} \left(\frac{CD}{\sqrt{n}}\right)^n \leq \mu^{-(1-\eta/2)n} \left(\frac{CD}{\sqrt{n}}\right)^n.$$

Finally, recalling that $N_1 \cdot N_2$ is less than the size of actual number of translates by at most $C^n D^2$, where C depends on K, K', λ, gives the desired conclusion. $\qquad\square$

We are now ready to prove Theorem 1.3; the proof is essentially identical to the proof in [12] except for one twist.

Proof of Theorem 1.3. By Lemma 2.4 applied to $A + M$ and $(A + M)^T$, it follows that there exist $\delta_{c_{2.4}}, \rho_{c_{2.4}}, c_{2.4} > 0$ depending only on η, K, and the sub-Gaussian norm of ξ such that except with probability $4e^{-c_{2.4}n}$, any left or right compressible vector has image with norm at least $c_{2.4}\sqrt{n}$. Then, by Lemma 2.5, it suffices to provide a uniform (in $k \in [n]$) bound on $\mathbb{P}(\text{dist}(X_k, H_k) \leq \epsilon)$, where X_k denotes the kth row of $A + M$, and H_k denotes the span of all the other rows of $A + M$. Henceforth, we will take $k = n$, noting that the argument for other values of k is exactly the same.

Let $v(\cdot)$ be a function mapping $(n-1)\times n$ matrices to an arbitrary unit vector in their right kernel. As before, let B denote the $(n-1)\times n$ matrix formed by the top $n-1$ rows of A, and let N denote the $(n-1)\times n$ matrix formed by the top $n-1$ rows of M. By Lemma 2.4, it follows that, except with probability at most $2e^{-c_{2.4}n}$, $v(B+N)$ must lie in $\mathrm{Incomp}_{\delta_{2.4},\rho_{2.4}}$. Therefore, by Lemma 2.7, there exist some $\gamma,\lambda>0$ (depending only on $\delta_{2.4},\rho_{2.4}$) such that $\mathrm{LCD}_{\alpha,\gamma}(v(B+N))\geq\lambda\sqrt{n}$ for any $\alpha>0$. We write $\alpha=\mu\sqrt{n}$, where $\mu>0$ will be chosen later; it is important to note that none of the previously chosen parameters depend on μ.

Let X_n denote the last row of $A+M$. Since $\mathrm{dist}(X_n,H_n)\geq|\langle v(B+N),X_n\rangle|$, it suffices to bound

$$\mathbb{P}[|\langle v(B+N),X_n\rangle|\leq\epsilon].$$

First, we consider, for dyadically chosen $D\in[\lambda\sqrt{n},2^{\chi n}]$ (where $\chi>0$ will be chosen at the end), the probability

$$\mathbb{P}[|\langle v(B+N),X_n\rangle|\leq\epsilon\wedge v(B+N)\in S_D].$$

The key observation here is that this probability is at most

$$\mathbb{P}[|\langle v(B+N),X_n\rangle|\leq\epsilon\wedge v(B+N)\in S_D'(K')]+\mathbb{P}[\|N\|\geq K'\sqrt{n}],$$

since for any $x\in S_D'(K')\backslash S_D$, on the event $\|N\|\leq K'\sqrt{n}$, we have $\|(B+N)x\|\geq 2K'\sqrt{n}-K'\sqrt{n}\neq 0$. Note that there exists $K'>0$ depending only on the sub-Gaussian norm of ξ for which the second term is exponentially small (cf. [12, Lemma 2.4]). We fix such a K'.

Now, we bound the first term as follows. Let $\mathcal{N}$ denote the net for $S_D'(K')$ coming from Lemma 2.14. Then,

$$\begin{aligned}
\mathbb{P}[|\langle v(B+N),X_n\rangle|\leq\epsilon\wedge v(B+N)\in S_D'(K')]\\
\leq\mathbb{P}[v(B+N)\in S_D'(K')]\\
\leq\sum_{w\in\mathcal{N}}\mathbb{P}[|(B+N)w|\leq\mu n/D]\\
\leq\mu^{-(1-\eta/2)n}D^2\left(\frac{C_{2.14}(\lambda,K,K')D}{\sqrt{n}}\right)^n\times\left(\frac{C\cdot C_{2.8}(\gamma)\mu\sqrt{n}}{D}\right)^{n-1},
\end{aligned}$$

which is exponentially small (here, C is an absolute constant), as long as $\mu>0$ is chosen to be sufficiently small, and then χ is chosen to be sufficiently small. In the last line, we have used Lemma 2.14, Lemma 2.8 and a standard tensorization argument (cf. [12, Lemma 2.2]).

Finally, we note that by Lemma 2.5, we have

$$\mathbb{P}[\langle v(B+N),A_n+M_n\rangle\leq\epsilon\wedge\mathrm{LCD}_{\mu\sqrt{n},\gamma}(v(B+N))\geq 2^{\chi n}]\leq C_{2.8}(\gamma)(\epsilon+2^{-\chi n}+2^{-c\mu^2 n}),$$

which completes the proof. $\qquad\square$

3. Proof of Theorem 1.4

Let A,R_n be defined as in the statement of Theorem 1.4. For convenience, we will study the least singular value of $A-R_n$, which has the same distribution as $A+R_n$. Let $K\geq 1$ be sufficiently large so that the event $\mathcal{E}_K=\{\|R_n\|\leq K\sqrt{n}\}$ satisfies

$$\mathbb{P}[\mathcal{E}_K]\geq 1-2^{-n}.$$

Finally, let $L\geq 2K\sqrt{n}$.

We first reduce the study of the smallest singular value to the study of anti-concentration of a very structured vector.

Lemma 3.1. *Let $u, w \in \mathbb{R}^{n-1}$ have i.i.d coordinates distributed as lazy Rademacher random variables. Then, we have*

$$\frac{1}{2}\mathbb{P}\left[\left|\sum_{i\geq 0} w^T(R_{n-1}/L)^i u\right| \leq L\epsilon \cap \|R_{n-1}\| \leq K\sqrt{n}\right] \leq \mathbb{P}[s_n(A - R_n) \leq \epsilon] + \mathbb{P}[\mathcal{E}_K^c].$$

Proof. Since we have the term $\mathbb{P}[\mathcal{E}_K^c]$ on the right hand side, we may henceforth restrict ourselves to the event $\|R_{n-1}\| \leq \|R_n\| \leq K\sqrt{n}$. On this event, note that for any $L \geq 2K\sqrt{n}$, the $(n-1) \times (n-1)$ random matrix $LI_{n-1} - R_{n-1}$ is invertible, and in fact, its inverse is expressible as a Neumann series, i.e.

$$(LI_{n-1} - R_{n-1})^{-1} = \sum_{i\geq 0} R_{n-1}^i/L^{i+1}.$$

The key is to decompose R_n as

$$R_n = \begin{pmatrix} R_{n-1} & u \\ w^T & r_{nn} \end{pmatrix};$$

note that the four block matrices appearing in the decomposition are independent of each other. Therefore, for

$$v = \begin{pmatrix} (LI_{n-1} - R_{n-1})^{-1}u \\ 1 \end{pmatrix},$$

we have trivially (from the last coordinate) that $\|v\|_2 \geq 1$, and hence,

$$
\begin{aligned}
s_n(A - R_n) &\leq \|(A - R_n)v\| \\
&\leq |w^T(L - R_{n-1})^{-1}u + r_{nn}| \\
&= |w^T(L - R_{n-1})^{-1}u|,
\end{aligned}
$$

where the final equality holds on the event $r_{nn} = 0$.

Finally, writing $(L - R_{n-1})^{-1}u = L^{-1}\sum_{i\geq 0}(R_{n-1}/L)^i u$, we have

$$
\begin{aligned}
\frac{1}{2}\mathbb{P}&\left[\left|\sum_{i\geq 0} w^T(R_{n-1}/L)^i u\right| \leq L\epsilon \cap \{\|R_{n-1}\| \leq K\sqrt{n}\}\right] \\
&= \mathbb{P}\left[\left|\sum_{i\geq 0} w^T(R_{n-1}/L)^i u\right| \leq L\epsilon \cap \{\|R_{n-1}\| \leq K\sqrt{n}\} \cap \{r_{nn} = 0\}\right] \\
&\leq \mathbb{P}\left[\left|\sum_{i\geq 0} w^T(R_{n-1}/L)^i u\right| \leq L\epsilon \cap \mathcal{E}_K \cap \{r_{nn} = 0\}\right] + \mathbb{P}[\mathcal{E}_K^c] \\
&\leq \mathbb{P}[s_n(A - R_n) \leq \epsilon] + \mathbb{P}[\mathcal{E}_K^c]. \qquad \square
\end{aligned}
$$

We are now in position to prove Theorem 1.4.

Proof of Theorem 1.4. Fix an integer $t \geq 1$, and let $C > 0$ be an absolute constant to be chosen later. Let

$$\mathcal{G}_K = \{Q \in \mathbb{R}^{(n-1)\times(n-1)} : \|Q\| \leq K\sqrt{n}\}.$$

Step 1: For independent R_{n-1}, u, w, consider the event

$$\mathcal{E} = \bigcap_{i=0}^{t-2}\{w^T R_{n-1}^i u = 0\} \cap \left\{\left|w^T\left(\sum_{i\geq t-1}(R_{n-1}/L)^i\right)u\right| \leq \frac{CK^t n^{t/2}}{L^{t-1}}\right\}.$$

12

Then, by Lemma 3.1 and the bound on $\mathbb{P}[\mathcal{E}_K]$, we have

$$\mathbb{P}\left[s_n(A - R_n) \leq C\left(\frac{K\sqrt{n}}{L}\right)^t\right]$$

$$\geq \frac{1}{2}\mathbb{P}\left[\left|\sum_{i\geq 0} w^T(R_{n-1}/L)^i u\right| \leq \frac{CK^{t-1}n^{t/2}}{L^{t-1}} \cap \{\|R_{n-1}\| \leq K\sqrt{n}\}\right] - 2^{-n}$$

$$\geq \frac{1}{2}\mathbb{P}[\mathcal{E} \cap \{R_{n-1} \in \mathcal{G}_K\}] - 2^{-n}.$$

Step 2: We write $u = u_1 - u_2$, where $u_1, u_2 \in \mathbb{R}^n$ are independent random vectors with i.i.d coordinates distributed as $\mathrm{Ber}(1/2) - 1/2$. For $j \in \{1, 2\}$, let $\mathcal{E}_j$ be the event

$$\mathcal{E}_j = \bigcap_{i=0}^{t-2}\{|w^T R_{n-1}^i u_j| \leq CK^i(\log 2t)n^{(i+1)/2}\} \cap \left\{\left|w^T\left(\sum_{i\geq t-1}(R_{n-1}/L)^i\right)u_j\right| \leq \frac{CK^{t-1}n^{t/2}}{2L^{t-1}}\right\}.$$

On the event $\{R_{n-1} \in \mathcal{G}_K\}$, we have $\|R_{n-1}^i\| \leq (K\sqrt{n})^i$, so that $\|R_{n-1}^i\|_{\mathrm{HS}} \leq K^i n^{(i+1)/2}$. Therefore, on this event, for any $L \geq 2K\sqrt{n}$, we have

$$\left\|\sum_{i\geq t-1}(R_{n-1}/L)^i\right\|_{\mathrm{HS}} \leq \frac{2K^{t-1}n^{t/2}}{L^{t-1}}.$$

Step 3: For independent random vectors u, w as above, and a fixed $(n-1) \times (n-1)$ matrix Q, consider the random quadratic polynomial $w^T Q u$. It follows from standard hypercontractive estimates (cf. [9, Theorem 10.24]), that there exists an absolute constant $c > 0$ such that for all $x \geq 0$,

$$\mathbb{P}[|w^T Q u| \geq x\|Q\|_{\mathrm{HS}}] \leq c^{-1}\exp(-cx).$$

From this, the Hilbert-Schmidt norm bounds in the previous step, and the union bound, it follows that for $j \in \{1, 2\}$, and for any $R_{n-1} \in \mathcal{G}_K$, we have

$$\mathbb{P}_{u_j, w}[\mathcal{E}_j(u_j, w, R_{n-1})] \geq \frac{1}{2},$$

provided that $C > 0$ is chosen sufficiently large.

Fix $R_{n-1} \in \mathcal{G}_K$. For $j \in \{1, 2\}$, let $W_j(R_{n-1})$ be the set of those vectors $w \in \{-1, 0, 1\}^{n-1}$ for which

$$\mathbb{P}_{u_j}[\mathcal{E}_j(u_j, w, R_{n-1})] \geq \frac{1}{4}.$$

Note that since u_1 and u_2 are identically distributed, $W_1(R_{n-1}) = W_2(R_{n-1}) =: W(R_{n-1})$. Then, by combining the conclusion of the previous step with averaging (reverse Markov's inequality), it follows that

$$\mathbb{P}[w \in W(R_{n-1})] \geq \frac{1}{4}.$$

Step 4: Fix $R_{n-1} \in \mathcal{G}_K$ and $w \in W(R_{n-1})$. Then, we have

$$\mathbb{P}_u[\mathcal{E}(u, w, R_{n-1})] \geq \mathbb{P}_{u_1, u_2}[\mathcal{E}(u_1 - u_2, w, R_{n-1}) \cap \mathcal{E}_1(u_1, w, R_{n-1}) \cap \mathcal{E}_2(u_2, w, R_{n-1})]$$

$$= \mathbb{P}_{u_1, u_2}[\mathcal{E}(u_1 - u_2, w, R_{n-1}) \mid \mathcal{E}_1(u_1, w, R_{n-1}) \cap \mathcal{E}_2(u_2, w, R_{n-1})] \times$$

$$\times \mathbb{P}_{u_1}[\mathcal{E}_1(u_1, w, R_{n-1})]\mathbb{P}_{u_2}[\mathcal{E}_2(u_2, w, R_{n-1})]$$

$$\geq \frac{1}{16}\mathbb{P}_{u_1, u_2}[\mathcal{E}(u_1 - u_2, w, R_{n-1}) \mid \mathcal{E}_1(u_1, w, R_{n-1}) \cap \mathcal{E}_2(u_2, w, R_{n-1})],$$

where the final inequality uses that $w \in W(R_{n-1})$.

Let us now bound from below the very last term. Conditioned on the event $\mathcal{E}_1(u_1, w, R_{n-1}) \cap \mathcal{E}_2(u_2, w, R_{n-1})$, we know that for $j \in \{1, 2\}$,

$$\left| w^T \left(\sum_{i \geq t-1} (R_{n-1}/L)^i \right) u_j \right| \leq \frac{CK^{t-1} n^{t/2}}{2L^{t-1}},$$

so that by the triangle inequality,

$$\left| w^T \left(\sum_{i \geq t-1} (R_{n-1}/L)^i \right) u \right| \leq \frac{CK^{t-1} n^{t/2}}{2L^{t-1}}.$$

Moreover, conditioned on the event $\mathcal{E}_1(u_1, w, R_{n-1}) \cap \mathcal{E}_2(u_2, w, R_{n-1})$, we know that the $(t-1)$-dimensional vectors

$$(w^T R_i^{n-1} u_j)_{0 \leq i \leq t-2}$$

are i.i.d for $j \in \{1, 2\}$, have each coordinate equal to a half-integer, and lie in a $(t-1)$-dimensional region with total number of points in $(\mathbb{Z}/2)^{(t-1)}$ at most

$$4^{t-1} C^{t-1} K^{\frac{(t-1)(t-2)}{2}} (\log 2t)^{t-1} n^{\frac{t(t-1)}{4}}.$$

In particular, by Cauchy-Schwarz, we see that these two $(t-1)$-dimensional vectors coincide (conditioned on $\mathcal{E}_1(u_1, w, R_{n-1}) \cap \mathcal{E}_2(u_2, w, R_{n-1})$) is at least

$$(4^{t-1} C^{t-1} K^{\frac{(t-1)(t-2)}{2}} (\log 2t)^{t-1} n^{\frac{t(t-1)}{4}})^{-1}.$$

But whenever this happens, we must also have

$$\bigcap_{i=0}^{t-2} \{ w^T R_{n-1}^i u = 0 \}.$$

To summarize, we have shown that for any $R_{n-1} \in \mathcal{G}_K$ and for any $w \in W(R_{n-1})$,

$$\mathbb{P}_u[\mathcal{E}(u, w, R_{n-1})] \geq 4^{-t-3} C^{1-t} K^{-\frac{(t-1)(t-2)}{2}} (\log 2t)^{1-t} n^{-\frac{t(t-1)}{4}}.$$

Step 5: The desired lower bound now follows easily. We have

$$
\begin{aligned}
\mathbb{P}_{u,w,R_{n-1}}[\mathcal{E} \cap \{R_{n-1} \in \mathcal{G}_K\}] &\geq \mathbb{P}_{u,w}[\mathcal{E}(u, w, R_{n-1}) | R_{n-1} \in \mathcal{G}_K](1 - 2^{-n}) \\
&\geq \mathbb{P}_{u,w}[\mathcal{E}(u, w, R_{n-1}) \cap \{w \in W(R_{n-1})\} | R_{n-1} \in \mathcal{G}_K](1 - 2^{-n}) \\
&\geq \mathbb{P}_u[\mathcal{E}(u, w, R_{n-1}) | \{w \in W(R_{n-1})\} \cap \{R_{n-1} \in \mathcal{G}_K\}] \times \\
&\quad \times \mathbb{P}_w[w \in W(R_{n-1}) | R_{n-1} \in \mathcal{G}_K](1 - 2^{-n}) \\
&\geq \frac{(1 - 2^{-n})}{4} \mathbb{P}_u[\mathcal{E}(u, w, R_{n-1}) | \{w \in W(R_{n-1})\} \cap \{R_{n-1} \in \mathcal{G}_K\}] \\
&\geq \frac{(1 - 2^{-n})}{4} 4^{-t-3} C^{1-t} K^{-\frac{(t-1)(t-2)}{2}} (\log 2t)^{1-t} n^{-\frac{t(t-1)}{4}},
\end{aligned}
$$

where the final inequality follows from the end of Step 4. Combining this with Step 1 completes the proof. $\square$

References

[1] Emmanuel Abbe, Jianqing Fan, Kaizheng Wang, Yiqiao Zhong, et al., *Entrywise eigenvector analysis of random matrices with low expected rank*, Annals of Statistics **48** (2020), 1452–1474.

[2] Zhidong Bai and Jack W Silverstein, *Spectral analysis of large dimensional random matrices*, vol. 20, Springer, 2010.

[3] Jess Banks, Archit Kulkarni, Satyaki Mukherjee, and Nikhil Srivastava, *Gaussian regularization of the pseudospectrum and Davies' conjecture*, arXiv:1906.11819.

[4] Jean Bourgain, *On a problem of Farrell and Vershynin in random matrix theory*, Geometric aspects of functional analysis, Springer, 2017, pp. 65–69.

[5] Alan Edelman, *Eigenvalues and condition numbers of random matrices*, SIAM J. Matrix Anal. Appl. **9** (1988), 543–560.

[6] Brendan Farrell and Roman Vershynin, *Smoothed analysis of symmetric random matrices with continuous distributions*, Proceedings of the American Mathematical Society **144** (2016), 2257–2261.

[7] Vishesh Jain, *Quantitative invertibility of random matrices: a combinatorial perspective*, arXiv:1908.11255.

[8] Galyna V Livshyts, Konstantin Tikhomirov, and Roman Vershynin, *The smallest singular value of inhomogeneous square random matrices*, arXiv:1909.04219.

[9] Ryan O'Donnell, *Analysis of Boolean functions*, Cambridge University Press, New York, 2014.

[10] Sean O'Rourke, Van Vu, and Ke Wang, *Random perturbation of low rank matrices: Improving classical bounds*, Linear Algebra and its Applications **540** (2018), 26–59.

[11] Mark Rudelson, *Recent developments in non-asymptotic theory of random matrices*, Modern aspects of random matrix theory, Proc. Sympos. Appl. Math., vol. 72, Amer. Math. Soc., Providence, RI, 2014, pp. 83–120.

[12] Mark Rudelson and Roman Vershynin, *The Littlewood-Offord problem and invertibility of random matrices*, Adv. Math. **218** (2008), 600–633.

[13] Mark Rudelson and Roman Vershynin, *Hanson-Wright inequality and sub-Gaussian concentration*, Electron. Commun. Probab. **18** (2013), no. 82, 9.

[14] Mark Rudelson and Roman Vershynin, *Small ball probabilities for linear images of high-dimensional distributions*, Int. Math. Res. Not. IMRN (2015), 9594–9617.

[15] Arvind Sankar, Daniel A Spielman, and Shang-Hua Teng, *Smoothed analysis of the condition numbers and growth factors of matrices*, SIAM Journal on Matrix Analysis and Applications **28** (2006), 446–476.

[16] Steve Smale, *On the efficiency of algorithms of analysis*, Bulletin (New Series) of The American Mathematical Society **13** (1985), 87–121.

[17] Daniel A Spielman and Shang-Hua Teng, *Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time*, Journal of the ACM (JACM) **51** (2004), 385–463.

[18] Terence Tao and Van Vu, *Random matrices: the circular law*, Communications in Contemporary Mathematics **10** (2008), 261–307.

[19] Terence Tao and Van Vu, *Smooth analysis of the condition number and the least singular value*, Math. Comp. **79** (2010), 2333–2352.

[20] Konstantin Tikhomirov, *Invertibility via distance for noncentered random matrices with continuous distributions*, Random Structures Algorithms **57** (2020), 526–562.

[21] Konstantin E Tikhomirov, *The smallest singular value of random rectangular matrices with no moment assumptions on entries*, Israel Journal of Mathematics **212** (2016), 289–314.

[22] Roman Vershynin, *High-dimensional probability*, Cambridge Series in Statistical and Probabilistic Mathematics, vol. 47, Cambridge University Press, Cambridge, 2018, An introduction with applications in data science, With a foreword by Sara van de Geer.

SIMONS INSTITUTE FOR THE THEORY OF COMPUTING, BERKELEY, CA 94720, USA
E-mail address: `visheshj@stanford.edu`

DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE, MA 02139, USA
E-mail address: `{asah,msawhney}@mit.edu`