

Smoothed Analysis of the Condition Number Under Low-Rank Perturbations

Rikhav Shah ✉

University of California at Berkeley, CA, USA

Sandeep Silwal ✉

Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

Let M be an arbitrary n by n matrix of rank $n - k$. We study the condition number of M plus a *low-rank* perturbation UV^T where U, V are n by k random Gaussian matrices. Under some necessary assumptions, it is shown that $M + UV^T$ is unlikely to have a large condition number. The main advantages of this kind of perturbation over the well-studied dense Gaussian perturbation, where every entry is independently perturbed, is the $O(nk)$ cost to store U, V and the $O(nk)$ increase in time complexity for performing the matrix-vector multiplication $(M + UV^T)x$. This improves the $\Omega(n^2)$ space and time complexity increase required by a dense perturbation, which is especially burdensome if M is originally sparse. Our results also extend to the case where U and V have rank larger than k and to symmetric and complex settings. We also give an application to linear systems solving and perform some numerical experiments. Lastly, barriers in applying low-rank noise to other problems studied in the smoothed analysis framework are discussed.

2012 ACM Subject Classification Mathematics of computing → Numerical analysis; Theory of computation → Randomness, geometry and discrete structures

Keywords and phrases Smoothed analysis, condition number, low rank noise

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2021.40

Category RANDOM

Related Version *Full Version*: <https://arxiv.org/abs/2009.01986>

Funding *Sandeep Silwal*: Research supported by the NSF Graduate Research Fellowship under Grant No. 1122374.

Acknowledgements We thank Sushruth Reddy and Samson Zhou for helpful conversations. We also thank Piotr Indyk and Arsen Vasilyan for helpful feedback on a draft of the paper.

1 Introduction

The smoothed analysis framework as introduced by Spielman and Teng aims to explain the performance of algorithms on real world inputs through a hybrid of worst-case and average case analysis [25]. In this framework, we are given an arbitrary input that is then perturbed randomly according to some specified noise model. We apply this framework to study the condition number of a matrix perturbed with low-rank Gaussian noise. The condition number is of interest since it influences the behavior of many algorithms in numerical linear algebra, both in theory and in practice.

To give context to our result, recall that the condition number of a $n \times n$ matrix M with singular values $s_1(M) \geq \dots \geq s_n(M)$ is defined as the ratio $s_1(M)/s_n(M)$. Generally, a condition number is “well behaved” if $s_1(M)/s_n(M) = n^{O(1)}$. It can be shown that under very mild and natural conditions, we have $s_1(M) \leq n^{O(1)}$. For instance, this readily follows from Proposition 6 if the entries are not too large compared to the size of M or if the entries have sufficient tail concentration far from the origin. Since our random variables are



© Rikhav Shah and Sandeep Silwal;

licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021).

Editors: Mary Wootters and Laura Sanità; Article No. 40; pp. 40:1–40:21



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Gaussians, this easily holds. Therefore, the bulk of the work lies in controlling the *smallest* singular value $s_n(M)$. Extending a result of Edelman [11], Sankar, Spielman and Teng showed the following result in [22]:

► **Theorem 1.** *There is a constant $C > 0$ such that the following holds. Let M be an arbitrary matrix and let N_n be a random matrix whose entries are iid Gaussian. Let $M_n = M + N_n$. Then for any $t > 0$,*

$$\mathbb{P}(s_n(M_n) \leq t) \leq Cn^{1/2}t.$$

Later, Tao and Vu generalized the above result where the entries of N_n are independent copies of a general class of random variables that have mean zero and bounded variance [32, 27].

1.1 Motivation for low-rank noise

The main drawback of these results is that *every* entry of M must be perturbed by independent noise. This means that if such a perturbation was carried out in practice, we would need to first draw n^2 random numbers and store them. This is more problematic if M is sparse to begin with and stored in a data structure utilized for sparse matrices. These observations lead us to ask if we can achieve well-conditioned matrices with **less randomness and less space**. Our results demonstrate the answer is yes by replacing the dense Gaussian ensemble N_n with a *low-rank* matrix.

To further motivate our work, we note that in the context of smoothed analysis, Theorem 1 is used to explain the phenomenon that matrices encountered in practice frequently have “well behaved” condition number. For instance, many matrices can arise out of empirical observations or measurements which can be subject to some inherent noise.

Similarly, low-rank noise is also natural and arises in many scenarios. Low-rank noise has been studied as a noise model in least squares [33], compressed sensing [8, 18, 10], and imaging [15, 26] to name a few applications. In addition, low-rank noise also arises in many applied sciences model, for instance, see the examples in [1] and references therein where examples are given for the eigenvalue problem $Mx = x + Ex$, for a low rank matrix E , arising in scientific modelling. Furthermore, one of the most frequent properties that matrices in data science possess is having low rank (see [6, 7, 30, 19] and references within). Thus for these matrices, the traditional smoothed analysis viewpoint of having a dense perturbation cannot apply due to their low rank requirement.

Hence, an additional motivation of our work is that studying low-rank noise is a natural step in smoothed analysis which we initiate.

1.2 Our results

As stated before, we replace the dense Gaussian perturbation N_n with a low-rank perturbation. Our main result is the following.

► **Theorem 2** (Theorem 13 simplified). *Let $1 \leq k \leq n/2$ and M be a matrix of rank $n - k$. Let U, V be $n \times k$ matrices with i.i.d. $\mathcal{N}(0, 1)$ entries. Then*

$$\mathbb{P}\left(s_n(M + UV^T) \leq \frac{\varepsilon}{nk} s_{n-k}(M)\right) \leq C\sqrt{\varepsilon} + \exp(-cn)$$

so long as $s_{n-k}(M) < n$ for absolute constants $C, c > 0$.

Theorem 2 roughly states that if we add a rank k random perturbation to a rank $n - k$ matrix, then the smallest k singular values of the matrix improve. The advantage of our approach is that the matrices U, V can be stored separately from M using $O(nk)$ space. This is especially useful in the case that M is sparse to begin with and is stored using a data structure optimized for sparse matrices. Furthermore, a matrix vector product operation $(M + UV^T)x$ can be computed in $\text{Time}(Mx) + O(nk)$ time where $\text{Time}(Mx)$ is the time required to compute Mx . For instance, when $k = O(1)$, the extra increase in space and time complexity is only $O(n)$. This is a significant improvement in both the space required to store a dense Gaussian random matrix G and computing $(M + G)x$ which are both $\Omega(n^2)$. We prove Theorem 2 in Section 3 and discuss the dependence on the terms $s_{n-k}(M)$ and $\sqrt{\varepsilon}$ which we show is unavoidable (see remarks 17, 18).

Theorem 2 can be generalized in a variety of ways. First, our result carries over to the case where we pick the columns of U, V to be from a rotationally invariant distribution, such as uniform vectors on the unit sphere. We show that our result also carries over to the case where M is symmetric and we pick $U = V$ to preserve symmetry.

It is natural to ask if a broader family of random variables can be used in Theorem 2. In Section 3.1 we show that our result *cannot hold* if we pick the entries of U, V to be from the Rademacher distribution. This is in contrast to the dense perturbation case where Gaussian random variables can be replaced with a wide variety of other distributions such as sub-Gaussian random variables (which include Rademachers). On the other hand, we can get an analogous statement to Theorem 13 if we allow for complex Gaussian perturbations.

► **Theorem 3** (Theorem 15 simplified). *Let $1 \leq k \leq n/2$ and M be a matrix of rank $n - k$. Let U, V be $n \times k$ complex matrices with real and imaginary parts in each entry drawn independently from $\mathcal{N}(0, 1/2)$. Then*

$$\mathbb{P}\left(s_n(M + UV^T) \leq \frac{\varepsilon}{nk} s_{n-k}(M)\right) \leq C\varepsilon + \exp(-cn)$$

so long as $s_{n-k}(M) < n$ for absolute constants $C, c > 0$.

A further natural question to consider is if the low-rank noise model can be studied in other problems in smoothed analysis. In Section B, we highlight the challenges that arise when applying low-rank random perturbations to other well studied problems in smoothed analysis such as the simplex method and k -means clustering. We show that current analysis methods that work for dense random perturbations for these problems do not carry over to the low-rank case due to the lack of independence.

Lastly, we note that Theorem 2 requires that if the input matrix M has rank $n - k$, then perturbation has rank exactly k . This condition can be relaxed in a couple of ways; first, if we add a perturbation of rank less than k then the matrix will be singular so there is nothing to study in this case. On the other hand, adding a rank $k' > k$ perturbation to a rank $n - k$ matrix can be thought of as adding a rank $(k' - k)$ perturbation to a full rank matrix since the original matrix plus a rank k perturbation will be full rank with probability 1. Then as explained further in Section 4, we can obtain the following result in this case.

► **Theorem 4** (Theorem 20 simplified). *Let M be a $n \times n$ real matrix with $\text{rank}(M) = n$, smallest singular value s_n , and $U, V \in \mathbb{R}^{n \times k}$ have independent $\mathcal{N}(0, 1)$ coordinates. Then for all $\varepsilon \in (0, 1)$,*

$$\mathbb{P}\left(s_n(M + UV^T) \leq \frac{\varepsilon}{\sqrt{n}}\right) \leq C \left(\sqrt{\varepsilon} \left(1 + \frac{2nk}{s_n}\right) + \frac{1}{(2nk)^{9/4}} + \exp(-cnk) \right).$$

The second way to circumvent Theorem 2 is with the use of Weyl’s perturbation inequality. To see how it applies, consider the case of $k = 1$. Decompose $M = s_n(M)\ell_n r_n^T + M'$ where ℓ_n, r_n are the left and right singular vectors associated with $s_n(M)$. Then we can view $M + uv^T$ as a random perturbation of M' (which has rank $n - 1$), plus matrix $s_n(M)\ell_n r_n^T$ whose operator norm is at most s_n . We can then apply Theorem 2 to M' to bound $s_n(M' + UV^T)$ in terms of $s_{n-1}(M') = s_{n-1}(M)$, and then incur an additional additive $s_n(M)$ error by Weyl’s inequality. Since our ideal use case is when $s_n(M)$ is already negligible, the final bound that we get is comparable to the bound from Theorem 2.

Finally in Section 3.2, we discuss an application of low-rank perturbations to solving large sparse linear systems and in Section A, we present numerical evidence for our low-rank error model.

► **Remark 5.** Note that Theorem 1 and the works of Tao and Vu in [32, 27] both prove a statement of the form $\mathbb{P}(s_n(M + E) \leq n^{-A}) \leq n^{-B}$ where E is the perturbation and A, B are parameters that depend on the random variables comprising the perturbation. Our statements in Theorem 2 is also of similar flavor since it shows that $\mathbb{P}(s_n(M + E) \leq s_{n-k}(M) n^{-A}) \leq n^{-B}$. Since the theorems of Sankar, Spielman, and Teng and Tao and Vu have found other applications in smoothed analysis, such as in the analysis of the simplex method and beyond, we envision that our theorem could also find similar applications. We discuss barriers in applying the low-rank noise model to other smoothed analysis problems in Section B.

1.3 Previous techniques and our approach

In summary, it is difficult to apply previous techniques in our case since we have *shared randomness* across different rows/columns of the matrix. In more detail, all of the previous techniques used to bound the singular values of a random matrix rely on the controlling the distance between a row to the span of the other rows. To see why this is relevant, imagine a singular matrix. In such a case, it is clear that there must exist a row that lies in the span of the other rows. Therefore, controlling the distance from a row to the span of the other rows gives control over the smallest singular value.

Controlling this geometric quantity boils down to understanding the dot product between a row and the normal vector of the hyperplane spanned by the other rows. Crucially if the rows are independent, we can treat the normal vector of the hyperplane as fixed so this question reduces to the well known Erdos-Littlewood-Offord anti-concentration inequality and its variants which are used in previous works such as [28, 32, 27].

To be more precise, lets consider a high level overview of the proof of Sankar, Spielman, and Teng’s Theorem 1. Fix a vector x and note that from the identity $s_n(M_n) = \|M_n^{-1}\|$, it suffices to give a tail bound on $\|M_n^{-1}x\|$. By applying an orthogonal rotation and using the rotational invariance of the Gaussian, we can say that

$$\|M_n^{-1}x\| = \|M_n^{-1}e_1\| = \|c_1\|$$

where e_1 is the first basis vector and c_1 is the first column of M_n^{-1} . From the equation $M_n \cdot M_n^{-1} = I$, it follows that $\|c_1\| = 1/|w^T r_1|$ where r_1 is the first row of M_n and w^T is the normal vector of the span of the rows r_2, \dots, r_n . Therefore, the proof reduces to understanding the dot product between a random vector r_1 , and another independent vector w . In the more general case of Tao and Vu [32, 27], more elaborate dot product estimates using the Erdos-Littlewood-Offord inequality are needed.

In our case, if we add a rank 1 perturbation to a matrix, randomness is shared across *all* rows. Therefore, we cannot reduce our problem to understanding the dot product between a random vector and another independent vector since fixing a normal hyperplane of a span

of a subset of rows automatically gives information about the rows not considered in the span due to the shared randomness. Hence, it is tricky to apply the spectrum of existing techniques in our case.

To overcome these barriers, we use a completely different method to prove Theorem 13. We first reduce our problem to adding noise to a diagonal matrix by using rotational invariance. Then we employ linear algebraic tools (rather than probabilistic tools), to get an “explicit” representation of the inverse of a matrix after adding rank k noise. After arriving at an explicit representation of the inverse, we are able to compute a probabilistic bound on the smallest singular value. Our proof crucially uses the fact that our low-rank perturbations have Gaussian entries whereas the proofs of the dense perturbations carry over in other distributional settings as well. This is not a flaw of our method since it is simply not possible to prove an analogue of our theorems even if the entries of the low-rank perturbations come from sub-Gaussian distributions. We elaborate this point further in Section 3.1.

For Theorem 20 where we add a rank k random noise to a rank n matrix, we carefully adapt the geometric ideas utilized in previous approaches as explained above. However to get around the shared randomness between rows, we have to perform some careful conditioning which allows us better control the behavior of the random normal vector w .

1.4 Why would the perturbed matrix even be full rank?

We briefly address the question of why we even expect low-rank perturbations to improve the condition number. Consider the case where D is a diagonal matrix of rank $n - 1$ and we add a random rank 1 Gaussian perturbation uv^T . Recall the matrix determinant lemma which states that

$$\det(D + uv^T) = \det(D) + v^T \text{adj}(D)u$$

where $\text{adj}(D)$ is the adjugate matrix of D . In our case, we can assume that the first $n - 1$ entries on the diagonal of D are given by $s_1(D), \dots, s_{n-1}(D)$ while the last entry is 0. Then, the adjugate matrix is the all zeros matrix except the bottom rightmost entry which is $s_1(D) \cdots s_{n-1}(D)$. Therefore,

$$\det(D) + v^T \text{adj}(D)u = (u_n v_n)(s_1(D) \cdots s_{n-1}(D))$$

which is non-zero with probability 1 since $u_n v_n \neq 0$ with probability 1. Thus, adding a random rank 1 perturbation results in D not being singular which motivates the question of studying the smallest singular value after a random rank 1 (and more generally low-rank) perturbation.

1.5 Related works

The smoothed analysis framework has been applied to a variety of problems, most notably in analyzing optimization problems such as k -means [2, 3], the perceptron algorithm [5], and the simplex method [25, 9]. In all of these results, the goal is to show that after an input instance of the problem is suitably perturbed, the algorithm or heuristic runs in polynomial time (the time may depend on the properties of the noise). For a survey of results, see [29, 24, 16] and references within. The analysis used tends to be very problem specific and also heavily dependent on the type of noise added which for a vast majority of cases are dense Gaussian noise.

Zero preserving noise

The work that is closest in spirit to our work is the zero preserving noise model studied by Spielman, and Teng. It was shown in [22] that if M is a *symmetric* matrix, then adding an independent Gaussian random variable x_{ij} to each entry M_{ij} such that $i \neq j$, $M_{ij} \neq 0$, and satisfying $x_{ij} = x_{ji}$ along with a Gaussian perturbation along the diagonal results in a “well behaved” condition number. However, the main drawback of this result is that it only holds for symmetric matrices and even in this case, a dense perturbation is required if M is dense to begin with.

Other works

There are works that use sparse Gaussian perturbations, i.e, their perturbation model is a Gaussian times a Bernoulli random variable with a small parameter. If the Bernoulli random variable has sufficiently small parameter, then with high probability, most of the entries in the perturbation will be zero. The downsides of these methods are that many random variables still need to be drawn and it is not clear if they can show the resulting matrix is well conditioned. For example, Theorem 3.6 in [20] only shows that the singular values of the resulting matrix are *separated* from each other, not that the matrix is well conditioned. In fact, the study of these types of random matrices where entries are sub-Gaussian random variables multiplied by independent Bernoulli variables is still lacking. For instance, the smallest singular value of such family of random matrices was only recently resolved in the highly technical paper of Basak and Rudelson [4].

1.6 Notation

We use capital letters as A, M to denote matrices and lower case letters such as x for vectors. For a vector x , the norm $\|x\|$ is always the Euclidean norm whereas for a matrix A , the norm $\|A\|$ always refers to the operator norm (the largest singular value). For a matrix A , let A_S denote the sub-matrix of A which includes the i th row of A if and only if $i \in S$. The relation $a \lesssim b$ denotes that a is less than or equal to b up to some fixed positive constant and similarly, $a \simeq b$ denotes that a and b are equal up to some fixed positive constant. Unless otherwise indicated, variables C, c, C_1, C_2, \dots denote positive constants.

2 Preliminaries

In this section we enumerate some useful results. First, we recall a classical estimate of the operator norm of a random matrix of Seginer [23]. The following proposition essentially shows that the top singular value of a random matrix is well behaved under mild assumptions. Alternatively, one can also bound the top singular value by the frobenius norm if the random variables populating the matrix have sufficient tail concentration.

► **Proposition 6.** *Let M be a random $n \times n$ matrix with entries m_{ij} . Then,*

$$\mathbb{E}\|M\| = O \left(\mathbb{E} \max_{1 \leq i \leq n} \sqrt{\sum_{j=1}^n m_{ij}^2} + \mathbb{E} \max_{1 \leq j \leq n} \sqrt{\sum_{i=1}^n m_{ij}^2} \right).$$

Next we establish tail bounds for the smallest and largest singular values of real and complex Gaussian matrices.

► **Lemma 7** (Theorem 1 reformulated.). Let $G \in \mathbb{R}^{k \times k}$ with all entries chosen i.i.d. from $\mathcal{N}(0, 1)$. Then

$$\mathbb{P}\left(s_k(G) \leq t_1/\sqrt{k}\right) < Ct_1.$$

for some absolute constant C .

► **Lemma 8** (Theorem 1.1 in [28]). Let $G \in \mathbb{C}^{k \times k}$ with all entries chosen with i.i.d. real and imaginary parts from $\mathcal{N}(0, 1/2)$. Then

$$\mathbb{P}\left(s_k(G) \leq t_1/\sqrt{k}\right) < t_1^2.$$

► **Lemma 9** (Proposition 2.3 in [21]). Let $G \in \mathbb{R}^{(n-k) \times k}$ for $k \leq n/2$ with all entries chosen i.i.d. from $\mathcal{N}(0, 1)$. Then

$$\mathbb{P}\left(s_1(G) \geq t_2\sqrt{n-k}\right) < C_1 e^{-C_2 t_2^2 n}.$$

for t_2 larger than some absolute constant, and C_1, C_2 absolute constants.

► **Lemma 10.** Let $G \in \mathbb{C}^{(n-k) \times k}$ for $k \leq n/2$ with all entries chosen with i.i.d. real and imaginary parts from $\mathcal{N}(0, 1/2)$. Then

$$\mathbb{P}\left(s_1(G) \geq t_2\sqrt{n-k}\right) < 2C_1 e^{-C_2 t_2^2 n}.$$

for t_2, C_1, C_2 as in Lemma 9.

Proof. Decompose $G = A + iB$ and bound $s_1(G) \leq s_1(A) + s_1(B)$. Then

$$\begin{aligned} \mathbb{P}\left(s_1(G) \geq t_2\sqrt{2(n-k)}\right) &\leq \mathbb{P}\left(s_1(A) + s_1(B) \geq t_2\sqrt{2(n-k)}\right) \\ &\leq \mathbb{P}\left(s_1(A) \geq \frac{t_2\sqrt{2(n-k)}}{2}\right) + \mathbb{P}\left(s_1(B) \geq \frac{t_2\sqrt{2(n-k)}}{2}\right) \\ &\leq \mathbb{P}\left(s_1(\sqrt{2}A) \geq t_2\sqrt{n-k}\right) + \mathbb{P}\left(s_1(\sqrt{2}B) \geq t_2\sqrt{n-k}\right) \\ &\leq 2C_1 e^{-C_2 t_2^2 n} \end{aligned}$$

where the last inequality follows by Lemma 9 since $\sqrt{2}A$ and $\sqrt{2}B$ have real i.i.d. $\mathcal{N}(0, 1)$ entries. ◀

The following lemma bounds the smallest singular value of a block matrix.

► **Lemma 11.** Let

$$M = \begin{bmatrix} A & B \\ C & D \end{bmatrix}$$

be an $n \times n$ matrix. Then

$$s_n(M)^{-1} \leq \|A^{-1}\| + \|(M/A)^{-1}\| (1 + \|A^{-1}B\|) (1 + \|CA^{-1}\|)$$

where $(M/A) = D - CA^{-1}B$ is the Schur complement of A .

Proof. We first use the Schur formula for the inverse of a block matrix:

$$M^{-1} = \begin{bmatrix} A^{-1} + A^{-1}B(M/A)^{-1}CA^{-1} & A^{-1}B(M/A)^{-1} \\ (M/A)^{-1}CA^{-1} & (M/A)^{-1} \end{bmatrix}.$$

The norm of M^{-1} is upper bounded by the sum of the norms of each of its blocks.

$$\begin{aligned} s_n(M)^{-1} = \|M^{-1}\| &\leq \|A^{-1}\| + \|A^{-1}B\| \|(M/A)^{-1}\| \|CA^{-1}\| \\ &\quad + \|A^{-1}B\| \|(M/A)^{-1}\| \\ &\quad + \|(M/A)^{-1}\| \|CA^{-1}\| \\ &\quad + \|(M/A)^{-1}\| \\ &= \|A^{-1}\| + \|(M/A)^{-1}\| (1 + \|A^{-1}B\|) (1 + \|CA^{-1}\|). \quad \blacktriangleleft \end{aligned}$$

Lastly, we recall that Gaussians are sufficiently anti-concentrated.

► **Proposition 12.** *Let $x \sim \mathcal{N}(0, 1)$. Then, $\mathbb{P}(|x| \leq \varepsilon) = \Theta(\varepsilon)$ for ε sufficiently small.*

3 Proof of main theorems

The goal of this section is to prove the following theorem and its complex and symmetric analogs.

► **Theorem 13.** *Let M be an arbitrary matrix of rank $n - k \geq n/2$. Let U, V be $n \times k$ matrices with i.i.d. $\mathcal{N}(0, 1)$ entries. Then*

$$\mathbb{P}\left(s_n(M + UV^T) \leq \frac{t_1^2}{k} \min\left(\frac{1}{2}, \frac{s_{n-k}(M)}{4t_2^2(n-k)}\right)\right) \leq C_1 t_1 + C_2 \exp(-C_3 t_2^2 n) \quad (1)$$

for $t_1 \leq C_4$ and $t_2 \geq C_5$ for some absolute constants $C_i, 1 \leq i \leq 5$.

Our strategy to prove Theorem 13 will reduce general M to the case of M nonnegative and diagonal, then express $s_n(M + UV^T)$ in terms of the singular values of M and certain sub-matrices of U and V , and finally apply tail bounds to said singular values. We start by proving a lemma that allows us to reduce to the case of M nonnegative and diagonal. As stated in Section 1.3, this is a completely different proof strategy than the one used in previous works.

► **Lemma 14.** *Let $D = \text{diag}(s_n(M), \dots, s_1(M))$. Let U, V be as in Theorem 13. Then the distributions of $s_n(M + UV^T)$ and of $s_n(D + UV^T)$ are identical.*

Proof. Let $LDR^T = M$ be the singular value decomposition of M . Then

$$M + UV^T = LDR^T + UV^T = L(D + L^T UV^T R)R^T.$$

Left- and right- multiplication by unitary matrices preserves singular values so

$$s_n(M + UV^T) = s_n(D + L^T UV^T R).$$

Finally, U and V are rotationally invariant, so $L^T U$ and $R^T V$ are distributed just as U and V are. ◀

Now we proceed to the main proof.

Proof of Theorem 13. For any matrix T , recall that T_S denotes the sub-matrix of T which includes the i th row of T if and only if $i \in S$. Lemma 14 shows that we may assume M is nonnegative and diagonal without loss of generality. We may write M and $M + UV^T$ in block form as

$$M = \begin{bmatrix} 0 & 0 \\ 0 & M' \end{bmatrix} \quad \text{and} \quad M + UV^T = \begin{bmatrix} U_{[k]}V_{[k]}^T & U_{[k]}V_{[n]\setminus[k]}^T \\ U_{[n]\setminus[k]}V_{[k]}^T & M' + U_{[n]\setminus[k]}V_{[n]\setminus[k]}^T \end{bmatrix}$$

where M' has no zeros on the diagonal. We can use Lemma 11 to upper bound $s_n(M + UV^T)$. The factor corresponding to the Schur complement is

$$\left\| \left(M' - U_{[n]\setminus[k]} \left(I - V_{[k]}^T (U_{[k]}V_{[k]}^T)^{-1} U_{[k]} \right) V_{[n]\setminus[k]}^T \right)^{-1} \right\| = \|M'^{-1}\| = s_{n-k}(M)^{-1}$$

since $I - V_{[k]}^T (U_{[k]}V_{[k]}^T)^{-1} U_{[k]} = 0$. This is one of the key insights of our proof. Then the resulting bound is

$$\begin{aligned} & s_n(M + UV^T)^{-1} \\ & \leq \frac{1}{s_k(U_{[k]})s_k(V_{[k]}^T)} + \frac{1}{s_{n-k}(M)} \left(1 + \|(U_{[k]}V_{[k]}^T)^{-1}U_{[k]}V_{[n]\setminus[k]}^T\| \right) \left(1 + \|U_{[n]\setminus[k]}V_{[k]}^T(U_{[k]}V_{[k]}^T)^{-1}\| \right) \\ & = \frac{1}{s_k(U_{[k]})s_k(V_{[k]}^T)} + \frac{1}{s_{n-k}(M)} \left(1 + \|V_{[k]}^{-1}V_{[n]\setminus[k]}^T\| \right) \left(1 + \|U_{[k]}^{-1}U_{[n]\setminus[k]}\| \right) \\ & \leq \frac{1}{s_k(U_{[k]})s_k(V_{[k]}^T)} + \frac{1}{s_{n-k}(M)} \left(1 + \|V_{[k]}^{-1}\| \|V_{[n]\setminus[k]}^T\| \right) \left(1 + \|U_{[k]}^{-1}\| \|U_{[n]\setminus[k]}\| \right) \\ & = \frac{1}{s_k(U_{[k]})s_k(V_{[k]}^T)} + \frac{1}{s_{n-k}(M)} \left(1 + \frac{s_1(V_{[n]\setminus[k]})}{s_k(V_{[k]})} \right) \left(1 + \frac{s_1(U_{[n]\setminus[k]})}{s_k(U_{[k]})} \right). \end{aligned}$$

Denote events

$$\begin{aligned} \mathcal{E}_1 &= \left(s_1(U_{[n]\setminus[k]}) \leq t_2\sqrt{n-k} \quad \text{and} \quad s_1(V_{[n]\setminus[k]}) \leq t_2\sqrt{n-k} \right), \\ \mathcal{E}_2 &= \left(s_k(U_{[k]}) \geq t_1/\sqrt{k} \quad \text{and} \quad s_k(V_{[k]}) \geq t_1/\sqrt{k} \right). \end{aligned}$$

Conditioning on \mathcal{E}_1 and \mathcal{E}_2 , the above bound becomes

$$s_n(M + UV^T)^{-1} \leq \frac{1}{s_{n-k}(M)} \left(1 + \frac{t_2}{t_1} \sqrt{(n-k)k} \right)^2 + \frac{k}{t_1^2}.$$

For sufficiently large n (specifically $n \geq 6\frac{t_1^2}{k t_2^2}$), this becomes

$$s_n(M + UV^T)^{-1} \leq \frac{k}{t_1^2} \left(\frac{2t_2^2(n-k)}{s_{n-k}(M)} + 1 \right) \leq \frac{2k}{t_1^2} \max \left(\frac{2t_2^2(n-k)}{s_{n-k}(M)}, 1 \right)$$

Taking the reciprocal of both sides yields

$$s_n(M + UV^T) \geq \frac{t_1^2}{2k} \min \left(\frac{s_{n-k}(M)}{2t_2^2(n-k)}, 1 \right)$$

The probability that this bound is violated is upper bounded by the probability that at least one of \mathcal{E}_1 or \mathcal{E}_2 fail. We may upper bound this quantity using the union bound:

$$\begin{aligned} \mathbb{P}(\neg\mathcal{E}_1 \vee \neg\mathcal{E}_2) &\leq \mathbb{P}(\neg\mathcal{E}_1) + \mathbb{P}(\neg\mathcal{E}_2) \\ &\leq \mathbb{P}(s_1(U_{[n]\setminus[k]}) \geq t_2\sqrt{n-k}) + \mathbb{P}(s_1(V_{[n]\setminus[k]}) \geq t_2\sqrt{n-k}) \\ &\quad + \mathbb{P}(s_k(U_{[k]}) \leq t_1/\sqrt{k}) + \mathbb{P}(s_k(V_{[k]}) \leq t_1/\sqrt{k}) \\ &\leq 2C_1t_1 + 2C_2e^{-C_3t_2^2n}. \end{aligned}$$

where the last step follows by applying Lemmas 7 and 9 twice each. The factors of 2 can be subsumed into the constants C_1 and C_2 giving the final result. \blacktriangleleft

40:10 Smoothed Analysis of the Condition Number Under Low-Rank Perturbations

► **Theorem 15.** *Let M, t_1, t_2, C_2, C_3 be as in theorem 13. Let U, V be $n \times k$ complex matrices with real and imaginary parts drawn independently from $\mathcal{N}(0, 1/2)$. Then*

$$\mathbb{P}\left(s_n(M + UV^T) \leq \frac{t_1^2}{k} \min\left(\frac{1}{2}, \frac{s_{n-k}(M)}{4t_2^2(n-k)}\right)\right) \leq 2t_1^2 + C_2 \exp(-C_3 t_2^2 n).$$

Note that the first term on the righthand side is $2t_1^2$ rather than $C_1 t_1$ as it was in Theorem 13.

Proof. The only place the proof differs from the proof of Theorem 13 is in the upper bound on $\mathbb{P}(-\mathcal{E}_1)$. Instead of $C_1 t_1$, it is simply t_1^2 by Lemma 8. ◀

► **Remark 16.** Theorems 13 and 15 hold when instead of sampling U and V independently, simply set $U = V$.

Proof. The proof follows almost exactly as before with only a single modification: In Lemma 14, the left- and right- singular vectors of symmetric matrices are the same so $L = R$ (so $L^T U = R^T V$). Optionally, one may note that events \mathcal{E}_1 and \mathcal{E}_2 are redundant, so one reduces the bound on $\mathbb{P}(-\mathcal{E}_1 \vee -\mathcal{E}_2)$ by a factor of 2. ◀

► **Remark 17.** Let us briefly mention why the term $s_{n-k}(M)$ is unavoidable in the statement of Theorem 13. For simplicity, consider $k = 1$ and suppose that M is of rank $n - 1$ and suppose its smallest nonzero singular value is equal to δ . After adding a rank 1 term uv^T to M , its rank is n with probability 1. However, if we consider the limit $\delta \rightarrow 0$, then $M + uv^T$ approaches a rank $n - 1$ matrix meaning $s_n(M + uv^T) \rightarrow 0$. Hence, any concentration bound such as (1) must depend on the term s_{n-1} .

► **Remark 18.** The term t_1 on the right hand side of (1) is also unavoidable. For simplicity, consider the case that $M \in \mathbb{R}^{n \times n}$ is a diagonal matrix and all the entries are non zero except the last diagonal entry and consider a rank 1 symmetric update. Now after a symmetric rank 1 update, the perturbed matrix is $M + gg^T$ where $g \in \mathbb{R}^n$. The smallest singular value of a symmetric matrix is equivalent to the smallest absolute eigenvalue. From the Raleigh quotient characterization of eigenvalues, we have that

$$\mathbb{P}(s_n(M + gg^T) \leq t^2) \geq \mathbb{P}(|e_n^T(M + gg^T)e_n| \leq t^2).$$

Using the fact that $Me_n = 0$, we have

$$\mathbb{P}(|e_n^T(M + gg^T)e_n| \leq t^2) = \mathbb{P}(z^2 \leq t^2)$$

where $z \in \mathbb{R}$ is a standard normal. Finally, $\mathbb{P}(z^2 \leq t^2) = \mathbb{P}(|z| \leq t) = \Theta(t)$ from Proposition 12 for t sufficiently small.

3.1 Sub-Gaussian perturbations

Just as Tao and Vu generalized Theorem 1 to the case where more general types of random perturbations beyond Gaussian are used, it is of interest to generalize Theorem 13 to the case where U, V are from a general family of distributions. A standard choice are mean zero sub-Gaussian distributions since they encompass well known distributions such as the standard Gaussian and ± 1 (Rademacher) random variables. Surprisingly, we show in this case that we cannot state a general statement like Theorem 13 unless extra assumptions about the fixed matrix M is made.

► **Lemma 19.** *Let $u, v \in \mathbb{R}^n$ be vectors with i.i.d. entries that are ± 1 (Rademacher) with equal probability. There exists a rank $n - 1$ matrix M such that with constant probability, $M + uv^T$ is singular.*

Proof. As in the proof of Theorem 13, let $M = LDR$ and we can say $s_n(M + uv^T) = s_n(D + (L^T u)(v^T R))$. In the case that u and v are Gaussian, rotational invariance implies that $L^T u$ and $v^T R$ are distributed as u, v respectively. However, this is no longer the case if u, v have entries coming from general sub-Gaussian distributions, such as ± 1 . Here, the properties of L, R can have substantial impact on $s_n(M + uv^T)$.

Suppose that the top left entry of D is 0. Then, if the first row of L is *sparse*, i.e. has $O(1)$ non zero entries, then it is possible that the first coordinate of $L^T u$, $(L^T u)_1$, is 0 with constant probability and hence the first row of $D + (L^T u)(v^T R)$ is all 0 which implies that $M + uv^T$ is still rank $n - 1$ with constant probability. ◀

Therefore, a general statement such as Theorem 13 in the case of sub-Gaussian distributions is impossible unless extra assumptions are made about the input matrix M . However, we note that in the $k = 1$ case, if we *assume* every row of L, R are *dense* (say have at least a constant fraction of non-zero entries), then the proof of Theorem 13 carries through in the ± 1 case since the two estimates we need (corresponding to the events \mathcal{E}_1 and \mathcal{E}_2 respectively) are the concentration of the norms of $L^T u, v^T R$ and each entry being anti-concentrated from 0 which follows from Erdos-Littlewood-Offord type results. It is not clear when such an assumption is natural.

3.2 Application to linear systems

We briefly highlight the importance of the condition number in solving systems of linear equations. If we are interested in solving the system $Ax = b$ where $A \in \mathbb{R}^{n \times n}$ then the condition number of A influences both the stability and runtime of linear systems solving. Much of this material is standard and can be found in [31].

Stability: If \tilde{x} denotes the result computed by numerical algorithms to the equation $Ax = b$ then it is known that the relative error quantity $\|x - \tilde{x}\|/\|x\|$ satisfies

$$\frac{\|x - \tilde{x}\|}{\|x\|} = O\left(\varepsilon_{\text{machine}} \cdot \frac{s_1(A)}{s_n(A)}\right)$$

where $\varepsilon_{\text{machine}}$ is the machine precision.

Runtime: One of the most widely used algorithms for solving systems of linear equations, especially large sparse ones that arise often in practice, is the conjugate gradient descent method. If the conjugate gradient descent method is run for k steps, then its convergence scales roughly as

$$\left(\frac{\sqrt{s_1(A)/s_n(A)} - 1}{\sqrt{s_1(A)/s_n(A)} + 1}\right)^k \approx \left(1 - \frac{2}{\sqrt{s_1(A)/s_n(A)}}\right)^k.$$

Therefore, a larger the condition number means more steps of the conjugate gradient descent method are required.

The usefulness of our low-rank error model is further supported by the conjugate gradient descent method. As mentioned previously, this iterative method is mainly used for large sparse systems. Thus, a low-rank perturbation that only requires additional linear space and incurs an additive linear increase in cost per iteration is desirable compared to a dense perturbation which makes the original problem infeasible for large systems.

4 Perturbation beyond rank k

In this section, we deal with the case that we add a rank k' perturbation to a rank $n - k$ matrix for $k' > k$. In such a case, we simply ignore a rank k portion of the noise and imagine that we are adding a rank $(k' - k)$ perturbation to a general full-rank matrix. This is valid since the original rank k matrix plus the rank k part of the noise will be full rank with probability 1. Our result then is the following.

► **Theorem 20.** *Let M be a $n \times n$ real matrix with $\text{rank}(M) = n > 10$ with smallest singular value s_n . Let $U, V \in \mathbb{R}^{n \times k}$ have independent $\mathcal{N}(0, 1)$ coordinates. Then,*

$$\mathbb{P}(s_n(M + UV^T) \leq 1/t) \leq C \left(\frac{\sqrt{n}}{x_2 \cdot t} \left(1 + \frac{x_1 \cdot x_3^{1/2} \cdot \sqrt{nk}}{s_n(M)} \right) + \exp(-x_1^2/4) + (2/\pi)^{k/2} x_2^k + \exp(-c x_3) \right) \quad (2)$$

for all $t > 0$, $x_1 \geq 3\sqrt{2 \log(2nk)}$, $x_3 \geq nk$, and $x_2 \leq 1$.

We obtain the following corollary under some natural parameter settings.

► **Corollary 21.** *Let M be a $n \times n$ real matrix with $\text{rank}(M) = n$, and $U, V \in \mathbb{R}^{n \times k}$ have independent $\mathcal{N}(0, 1)$ coordinates. Then for all $\varepsilon \in (0, 1)$,*

$$\mathbb{P}(s_n(M + UV^T) \leq \varepsilon/\sqrt{n}) \leq C \left(\sqrt{\varepsilon} \left(1 + \frac{3nk\sqrt{\log(2nk)}}{s_n(M)} \right) + \frac{1}{(2nk)^{9/4}} + \exp(-cnk) \right). \quad (3)$$

Proof. Set $t = \sqrt{n}/\varepsilon$, $x_1 = 3\sqrt{\log(2nk)}$, $x_2 = \sqrt{\varepsilon}$, $x_3 = nk$. ◀

By setting ε appropriately, we recover the common “theme” of $\mathbb{P}(s_n(M + UV^T) \leq n^{-A}) \leq n^{-B}$ as in the case of Theorem 1 and the works of Tao and Vu [32, 27].

We now proceed to prove Theorem 20. Denote $A = M + UV^T$ and note that $\text{rank}(A) = n$ with probability 1 so A^{-1} exists. We observe that (2) reduces to bounding $\mathbb{P}(\|A^{-1}\| \geq t)$. Our proof is adapted from the proof of Theorem 1. However, we need to perform a careful conditioning argument to prove the most important part of the argument which is presented in Lemma 23.

We begin by handling the case of a single vector.

► **Lemma 22.** *For any unit vector y , we have*

$$\mathbb{P}(\|A^{-1}y\|_2 \geq t) \leq C \left(\frac{1}{x_2 \cdot t} \left(1 + \frac{x_1 \cdot x_3^{1/2} \cdot \sqrt{nk}}{s_n(M)} \right) + \exp(-x_1^2/4) + (2/\pi)^{k/2} x_2^k + \exp(-c x_3) \right)$$

for all $t > 0$, $x_1 \geq 2\sqrt{\log n}$, $x_3 \geq nk$, and $x_2 \leq 1$.

Proof. Let Q be a rotation that takes y to e_n and denote QA as A' . Then,

$$\|A^{-1}y\|_2 = \|A^{-1}Q^T e_n\|_2 = \|(QA)^{-1}e_n\|_2 = \|c_n\|_2.$$

where c_n be the last column of $(QA)^{-1}$. From the identity $A'A'^{-1} = I$, we have that c_n is orthogonal to the first $n - 1$ rows of A' and has dot product 1 with the last row of A' . Hence,

$$\|c_n\|_2 = \frac{1}{|\langle w, r^n \rangle|}$$

where r^i is the i th row of A' and w is the unique unit vector orthogonal to the span of $\{r^1, \dots, r^{n-1}\}$ (up to to sign). This means that

$$\mathbb{P}(\|A^{-1}y\|_2 \geq t) = \mathbb{P}(\|c_n\|_2 \geq t) = \mathbb{P}(|\langle w, r^n \rangle| \leq 1/t).$$

The last row r^n of A' is the sum of the last rows of QM and QUV^T . Note that the inner product of w with the last row of QM is some fixed parameter; denote it r . Then QU is distributed as U by the rotational invariance of the normal distribution, so the last row of QUV^T is distributed as Vu^n where $u^n \in \mathbb{R}^k$ is a vector of independent Gaussians. Therefore, $\langle w, r^n \rangle$ is distributed as $\langle w, Vu^n \rangle + r$, so it suffices to bound the Levy concentration of $\langle w, Vu^n \rangle$. Specifically, we want to show that

$$\sup_{r \in \mathbb{R}} \mathbb{P}(|\langle w, Vu^n \rangle + r| < 1/t) \leq C \left(\frac{1}{x_2 \cdot t} \left(1 + \frac{x_1 \cdot x_3^{1/2} \cdot \sqrt{nk}}{s_n(M)} \right) + \exp(-x_1^2/4) + (2/\pi)^{k/2} x_2^k + \exp(-c x_3) \right)$$

where the probability is over the realization of u^n and V . This readily follows from Lemma 23. \blacktriangleleft

Proof of Theorem 20. Let s be a unit vector chosen uniformly at random from \mathcal{S}^{n-1} . By Lemma 22, we have

$$\mathbb{P}_{A,s}(\|A^{-1}s\|_2 \geq t/\sqrt{n}) \leq C \left(\frac{1}{x_2 \cdot t/\sqrt{n}} \left(1 + \frac{x_1 \cdot x_3^{1/2} \cdot \sqrt{nk}}{s_n(M)} \right) + \exp(-x_1^2/4) + (2/\pi)^{k/2} x_2^k + \exp(-c x_3) \right).$$

Now with probability 1, there exists a unique y such that $\|A^{-1}y\|_2 = \|A^{-1}\|$. From Lemma 24, we see that

$$\|A^{-1}s\|_2 \geq \|A^{-1}(y^T s)y\|_2 = |y^T s| \|A^{-1}\|.$$

Therefore we have

$$\begin{aligned} \mathbb{P}_{A,s}(\|A^{-1}s\|_2 \geq t/\sqrt{n}) &\geq \mathbb{P}_{A,s}(\|A^{-1}\| \geq t \text{ and } |s^T y| \geq 1/\sqrt{n}) \\ &= \mathbb{P}(\|A^{-1}\| \geq t) \cdot \mathbb{P}_{A,s}(|s^T y| \geq 1/\sqrt{n} \mid \|A^{-1}\| \geq t). \end{aligned}$$

By the rotational invariance of s , we have that

$$\mathbb{P}(\|A^{-1}\| \geq t) \cdot \mathbb{P}_{A,s}(|s^T y| \geq 1/\sqrt{n} \mid \|A^{-1}\| \geq t) = \mathbb{P}(\|A^{-1}\| \geq t) \cdot \mathbb{P}(|s^T e_1| \geq 1/\sqrt{n}).$$

From Lemma 25, we have that

$$\mathbb{P}(|s^T e_1| \geq 1/\sqrt{n}) \geq \mathbb{P}(|Z| \geq 1) - 1/n$$

where $Z \sim \mathcal{N}(0, 1)$. Altogether, it follows that

$$\begin{aligned} \mathbb{P}(\|A^{-1}\| \geq t) &\leq \frac{\mathbb{P}_{A,s}(\|A^{-1}s\|_2 \geq t/\sqrt{n})}{\mathbb{P}(|Z| \geq 1) - 1/n} \\ &\leq C \left(\frac{1}{x_2 \cdot t/\sqrt{n}} \left(1 + \frac{x_1 \cdot x_3^{1/2} \cdot \sqrt{nk}}{s_n(M)} \right) + \exp(-x_1^2/4) + (2/\pi)^{k/2} x_2^k + \exp(-c x_3) \right) \end{aligned}$$

for $n > 10$ as desired. \blacktriangleleft

► Lemma 23. Let $U, V \in \mathbb{R}^{n \times k}$ have independent $\mathcal{N}(0, 1)$ coordinates. Let $M \in \mathbb{R}^{n \times n}$ be a matrix with singular values $s_1 \geq \dots \geq s_n$ and let w be a vector perpendicular to the first $n-1$ rows of $M + UV^T$. Then for $x_1 \geq 3\sqrt{\log(2nk)}$, $x_3 \geq nk$, and $x_2 \leq 1$, and all $t > 0$,

$$\sup_{r \in \mathbb{R}} \mathbb{P}(|\langle w, Vu^n \rangle - r| < 1/t) \leq C \left(\frac{1}{x_2 \cdot t} \left(1 + \frac{x_1 \cdot x_3^{1/2} \cdot \sqrt{nk}}{s_n(M)} \right) + \exp(-x_1^2/4) + (2/\pi)^{k/2} x_2^k + \exp(-c x_3) \right)$$

for some $C, c > 0$ where u^n is the last row of U .

40:14 Smoothed Analysis of the Condition Number Under Low-Rank Perturbations

Proof. Let m^1, \dots, m^n be the rows of M and let u^1, \dots, u^n be the rows of U . Then the rows of $A = M + UV^T$ are given by $m^i + Vu^i$. Let y be the unit vector orthogonal to the span of $\{m^1, \dots, m^{n-1}\}$. Consider the following three events:

- $\mathcal{E}_1 =$ event that every entry of U is at most x_1 in absolute value,
- $\mathcal{E}_2 =$ event that $\|V^T y\|$ is at least x_2 ,
- $\mathcal{E}_3 =$ event that $\|V\|^2$ is at most x_3 ,

for $x_1 \geq 3\sqrt{\log(2nk)}, x_3 \geq nk$, and $x_2 \leq 1$. Denote $\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2 \cup \mathcal{E}_3$. We now show each of these occurs with high probability. By a standard concentration bound, the maximum of nk i.i.d. standard Gaussians is strongly concentrated around $\sqrt{2\log(2nk)}$. In particular, \mathcal{E}_1 happens with probability at least $1 - 2\exp(-x_1^2/4)$. Next, each coordinate of $V^T y$ is distributed as $\mathcal{N}(0, 1)$, and $\|V^T y\| \geq \|V^T y\|_\infty$, so we may upper bound $\mathbb{P}(\mathcal{E}_2^c)$ by $\mathbb{P}(|g| < x_2)^k$ where g is $\mathcal{N}(0, 1)$. This means \mathcal{E}_2 occurs with probability at least $1 - (2/\pi)^{k/2} x_2^k$. Lastly, the event \mathcal{E}_3 happens with probability at least $1 - \exp(-\Omega(x_3))$ by Lemma 9.

Now fix some realization of V and U such that \mathcal{E} occurs. Suppose for some parameter $z < \frac{\sqrt{3}}{2} \frac{s_n}{x_1 \cdot \sqrt{nk}}$ that we have $\|V^T w\| < z$. We will find a statement which this assumption implies, then take the contrapositive to obtain a lower bound on $\|V^T w\|$. From definition of w , we know

$$\langle w, m^i \rangle + \langle w, Vu^i \rangle = 0$$

for $1 \leq i \leq n-1$. We may apply Cauchy-Schwarz to $\langle w, Vu^i \rangle = \langle V^T w, u^i \rangle$ and use event \mathcal{E}_1 to bound $\|u^i\|$ and obtain

$$|\langle w, m^i \rangle| \leq \|V^T w\| \cdot \|u^i\| \leq z \cdot x_1 \cdot \sqrt{k}.$$

Decompose $w = w^\parallel + w^\perp$ where w^\parallel is in the span of m^1, \dots, m^{n-1} and w^\perp is in the orthogonal complement. Write $w^\parallel = \sum_{i=1}^{n-1} \alpha_i m^i$ for some coefficients α_i . Then we have

$$\begin{aligned} \|w^\parallel\|^2 &= |\langle w^\parallel, w \rangle| \leq \sum_{i=1}^{n-1} |\alpha_i| \cdot |\langle m^i, w \rangle| \\ &\leq \|\alpha\|_1 \cdot z \cdot x_1 \cdot \sqrt{k} \\ &\leq \|\alpha\| \cdot z \cdot x_1 \cdot \sqrt{nk} \\ &\leq \|w^\parallel\| \frac{z \cdot x_1 \cdot \sqrt{nk}}{s_n} \end{aligned}$$

where α is the vector of α_i 's, and the last step follows since M is non-singular, making α the unique solution to $M^T \alpha = w^\parallel$. Now, note that y and w^\perp are parallel, so

$$\|w^\perp\| \|V^T y\| = \|V^T w^\perp\| \leq \|V^T w\| + \|V^T w^\parallel\| \leq z + \|V\| \|w^\parallel\| \leq z \left(1 + \frac{x_1 \cdot x_3^{1/2} \cdot \sqrt{nk}}{s_n} \right)$$

where the last step follows from \mathcal{E}_3 . From $z < \frac{\sqrt{3}}{2} \frac{s_n}{x_1 \cdot \sqrt{nk}}$, we have

$$\|w^\perp\| = \sqrt{1 - \|w^\parallel\|^2} = \sqrt{1 - \left(\frac{z \cdot x_1 \cdot \sqrt{nk}}{s_n} \right)^2} > \frac{1}{2}.$$

Moving $\|w^\perp\|$ to the right hand side and using \mathcal{E}_2 and the above bound, we arrive at

$$x_2 \leq \|V^T y\| < 2z \left(1 + x_1 \cdot x_3^{1/2} \cdot \sqrt{nk}/s_n \right).$$

We have thus established the syllogism

$$\|V^T w\| \leq z \leq \frac{\sqrt{3}}{2} \frac{s_n}{x_1 \cdot \sqrt{nk}} \implies \frac{1}{2} \frac{x_2 s_n}{s_n + x_1 \cdot x_3^{1/2} \cdot \sqrt{nk}} < z.$$

By taking the contrapositive and setting $z = \frac{1}{2} \frac{x_2 s_n}{s_n + x_1 \cdot x_3^{1/2} \cdot \sqrt{nk}}$, we see that one of the inequalities on the left must fail. It isn't the second one since $x_2 < \sqrt{3}$, $x_3^{1/2} > 1$, and $s_n > 0$. We therefore conclude our selection of z lower bounds $\|V^T w\|$.

This leads to sufficient anti-concentration. For any fixed vector x , the distribution of $\langle x, u^n \rangle$ is the same as $\mathcal{N}(0, \|x\|^2)$. So, $\langle V^T w, u^n \rangle$ for random V is a mixture of Gaussians, each of which have variance at least z^2 when we condition on \mathcal{E} . The Levy concentration is thus easily bounded by

$$\sup_{r \in \mathbb{R}} \mathbb{P} \left(|\langle w, V u^n \rangle + r| \leq \frac{1}{t} \mid \mathcal{E} \right) \leq \frac{\sqrt{2/\pi}}{z \cdot t}.$$

The desired bound then follows by incorporating a probability bound on the complement of \mathcal{E} . \blacktriangleleft

The following lemmas follow easily from basic properties of the SVD and Gaussian random variables so we omit their proof.

► **Lemma 24.** Consider a $n \times n$ matrix M and $u \in \mathcal{S}^{n-1}$ such that $\|M\| = \|Mu\|_2$. Then for every $v \in \mathbb{R}^n$, we have

$$\|Mv\|_2 \geq |u^T v| \|M\|.$$

► **Lemma 25.** Let x be a uniformly random vector in \mathcal{S}^{n-1} and $Z \sim \mathcal{N}(0, 1)$. Then for every $c > 0$,

$$\mathbb{P} \left(|x^T e_1| \geq \sqrt{\frac{c}{n}} \right) \geq \mathbb{P}(|Z| \geq \sqrt{c}) - \frac{1}{n}.$$

References

- 1 Thomas J Anastasio, Andrea K Barreiro, and Jared C Bronski. A geometric method for eigenvalue problems with low-rank perturbations. *The Royal Society Publishing.*, 4, September 2017. doi:10.1098/rsos.170390.
- 2 D. Arthur and S. Vassilvitskii. Worst-case and smoothed analysis of the icp algorithm, with an application to the k-means method. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 153–164, 2006.
- 3 David Arthur, Bodo Manthey, and Heiko Röglin. Smoothed analysis of the k-means method. *J. ACM*, 58(5), 2011. doi:10.1145/2027216.2027217.
- 4 Anirban Basak and Mark Rudelson. Invertibility of sparse non-hermitian matrices. *Advances in Mathematics*, 310:426–483, 2017. doi:10.1016/j.aim.2017.02.009.
- 5 Avrim Blum and John Dunagan. Smoothed analysis of the perceptron algorithm for linear programming. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '02, page 905–914, USA, 2002. Society for Industrial and Applied Mathematics.
- 6 E. J. Candes and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010. doi:10.1109/JPROC.2009.2035722.
- 7 E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006. doi:10.1109/TIT.2005.862083.

- 8 Zhuo Chen and D. Ellis. Speech enhancement by sparse, low-rank, and dictionary spectrogram decomposition. *2013 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, pages 1–4, 2013.
- 9 Daniel Dadush and Sophie Huiberts. A friendly smoothed analysis of the simplex method. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, page 390–403, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3188745.3188826.
- 10 Minh Dao, Yuanming Suo, Sang Chin, and Trac Tran. Structured sparse representation with low-rank interference. *Conference Record - Asilomar Conference on Signals, Systems and Computers*, 2015:106–110, April 2015. doi:10.1109/ACSSC.2014.7094407.
- 11 Alan Edelman. Eigenvalues and condition numbers of random matrices. *SIAM Journal on Matrix Analysis and Applications*, 9(4):543–560, 1988. doi:10.1137/0609045.
- 12 Friedrich Eisenbrand and Santosh S. Vempala. Geometric random edge. *Math. Program.*, 164(1-2):325–339, 2017. doi:10.1007/s10107-016-1089-0.
- 13 Noam D. Elkies (<https://mathoverflow.net/users/14830/noam-d-elkies>). Smallest non-zero eigenvalue of a (0,1) matrix. MathOverflow. URL:<https://mathoverflow.net/q/157554> (version: 2017-04-13). arXiv:<https://mathoverflow.net/q/157554>.
- 14 V. Klee, G.J. Minty, and WASHINGTON UNIV SEATTLE Dept. of MATHEMATICS. *HOW GOOD IS THE SIMPLEX ALGORITHM*. Defense Technical Information Center, 1970. URL: <https://books.google.com/books?id=R8430AAACAAJ>.
- 15 X. Liu, M. Tanaka, and M. Okutomi. Practical signal-dependent noise parameter estimation from a single noisy image. *IEEE Transactions on Image Processing*, 23(10):4361–4371, 2014. doi:10.1109/TIP.2014.2347204.
- 16 Bodo Manthey and Heiko Raglin. Smoothed analysis: Analysis of algorithms beyond worst case. *it - Information Technology*, 53(6):280–286, 2011. doi:10.1524/itit.2011.0654.
- 17 Bodo Manthey and Rianne Veenstra. Smoothed analysis of the 2-opt heuristic for the tsp: Polynomial bounds for gaussian noise. In Leizhen Cai, Siu-Wing Cheng, and Tak-Wah Lam, editors, *Algorithms and Computation*, pages 579–589, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- 18 Morteza Mardani, Gonzalo Mateos, and G.B. Giannakis. Recovery of low-rank plus compressed sparse matrices with application to unveiling traffic anomalies. *IEEE Transactions on Information Theory*, 59, April 2012.
- 19 Sean O’Rourke, Van Vu, and Ke Wang. Random perturbation of low rank matrices: Improving classical bounds. *Linear Algebra and its Applications*, 540:26–59, 2018. doi:10.1016/j.laa.2017.11.014.
- 20 Richard Peng and Santosh Vempala. Solving sparse linear systems faster than matrix multiplication. In *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’21, page 504–521, USA, 2021. Society for Industrial and Applied Mathematics.
- 21 Mark Rudelson and Roman Vershynin. Smallest singular value of a random rectangular matrix. *Communications on Pure and Applied Mathematics*, 62(12):1707–1739, 2009. doi:10.1002/cpa.20294.
- 22 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(2):446–476, 2006. doi:10.1137/S0895479803436202.
- 23 Yoav Seginer. The expected norm of random matrices. *Combinatorics, Probability and Computing*, 9:149–166, March 2000. doi:10.1017/S096354830000420X.
- 24 Daniel A. Spielman and Shang hua Teng. Smoothed analysis: an attempt to explain the behavior of algorithms in practice. *COMMUN. ACM*, 2009.
- 25 Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *J. ACM*, 51(3):385–463, 2004. doi:10.1145/990308.990310.

- 26 Moeller Steen, Sebastian Weingartner, and Mehmet Akcakaya. Multi-scale locally low-rank noise reduction for high-resolution dynamic quantitative cardiac mri. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society.*, 2017:1473–1476, July 2017. doi:10.1109/EMBC.2017.8037113.
- 27 Terence Tao and Van Vu. Smooth analysis of the condition number and the least singular value. *Mathematics of Computation*, 79(272):2333–2352, 2010. URL: <http://www.jstor.org/stable/20779147>.
- 28 Vu Tau. Random matrices: the distribution of the smallest singular values. *Geometric and Functional Analysis*, 20:260–297, March 2010. doi:10.1007/s00039-010-0057-8.
- 29 Shang-Hua Teng. Smoothed analysis of algorithms and heuristics. In Lusheng Wang, editor, *Computing and Combinatorics*, pages 10–11, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- 30 C Tomasi and T. Kanade. Shape and motion from image streams under orthography: a factorization method. *Int Journal of Computer Vision*, 9:137–154, 1992. doi:10.1007/BF00129684.
- 31 Lloyd N. Trefethen and David Ill. Bau. *Numerical linear algebra*. SIAM Society for Industrial and Applied Mathematics, 2000.
- 32 Van H. Vu and Terence Tao. The condition number of a randomly perturbed matrix. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '07, page 248–255, New York, NY, USA, 2007. Association for Computing Machinery. doi:10.1145/1250790.1250828.
- 33 Karl Werner and Magnus Jansson. Parameter estimation for reduced-rank multivariate linear regressions in the presence of correlated noise. In *Conference Record of the Asilomar Conference on Signals, Systems and Computers*, volume 2, pages 2101–2105 Vol.2, December 2003.
- 34 Wikipedia contributors. Product distribution — Wikipedia, the free encyclopedia, 2020. [Online; accessed 18-August-2020]. URL: https://en.wikipedia.org/w/index.php?title=Product_distribution&oldid=970273508.

A Numerical experiments

In this section, we numerically demonstrate our theoretical results by giving an example of a sparse family of n by n matrices that are “poorly” conditioned and whose condition number improves significantly after adding a random Gaussian rank 1 perturbation. We show that this perturbation results in an improvement comparable to what is achieved after adding a dense Gaussian matrix while maintaining a low time complexity for matrix vector product operations.

Our family of n by n matrices will be constructed as follows: M_n will have ones on the anti-diagonal and the first and third off-diagonals above the anti-diagonal. For example, M_7 is displayed below.

$$\begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

It is shown in [13] that M_n is ill-conditioned by showing that the magnitude of the smallest eigenvalue of M_n is of the order $O(n/C^n)$ where $C \approx 1.47$ which implies that the smallest singular value of M_n is also at most $O(n/C^n)$. The second smallest eigenvalue on the other hand is a constant.

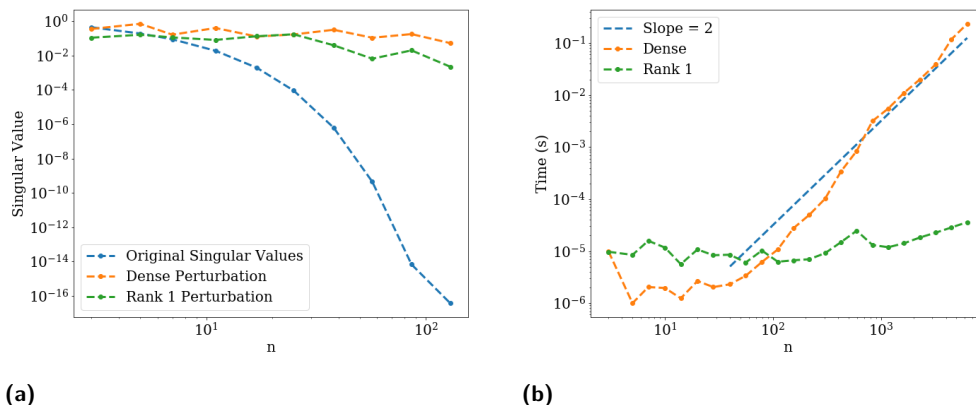


Figure 1 (a) Smallest singular values of the original matrix compared against dense and rank 1 perturbations. (b) Time taken to perform a matrix vector product after a dense perturbation and a rank 1 perturbation. The cost for the dense perturbation has a quadratic scaling (slope = 2).

In Figure 1 (a), we show the smallest singular value of M_n for a range of n along with the smallest singular values after a dense and rank 1 perturbation. As we can see in the log-log plot, the original values are decaying exponentially while the smallest singular value after the rank 1 perturbation is within a few orders of magnitude of the corresponding value after a dense perturbation. In Figure 1 (b), we show the time taken to perform a matrix vector product after a dense and a rank 1 perturbation. For this task, we used the popular numerical libraries NumPy and SciPy. Since M_n is sparse, it can be represented in a special sparse format to speed up computations. In the case of a rank 1 perturbation, we only need to store two additional vectors and a matrix vector product $(M + uv^T)x$ can be performed as

$$(M_n + uv^T)x = M_n x + (v^T x) \cdot u.$$

However, in the case of a dense perturbation, we need to store a dense matrix G and perform the matrix vector product operation with a vector and a dense matrix resulting in a much slower operation than in the rank 1 case. Indeed, note that the slope of the “Dense” curve in Figure 1 (b) is close to 2 signifying a quadratic increase in time. Overall, we see that in this case, a rank 1 update results in a comparable improvement of the condition number of M_n while greatly improving the cost to perform a fundamental matrix operation.

B Low-rank noise model for other problems in smoothed analysis

In this section, we outline some of the challenges that arise when applying the rank 1 noise model in other popular problems studied in smoothed analysis. While not a comprehensive survey of all problems, we focus on two of the most studied applications of this framework outside of the condition number. These are the simplex method and k -means. For these problems, the standard noise model is the dense one where every entry of the input matrix or input set of points respectively, is independently perturbed by a random Gaussian. We highlight some of the challenges that arise when trying to carry out existing proof techniques for these problems using rank 1 noise. This ultimately shows that new ideas are required to bypass the lack of independence as we did for the condition number.

B.1 Simplex method

The simplex method is one of the most famous applications of the smoothed analysis framework. The goal is to solve a linear program of the form $\max c^T x$ subject to $Ax \leq b$ using the simplex method where the entries of $A \in \mathbb{R}^{m \times n}$ have been perturbed by random noise. Recall that the simplex method operates by moving among the vertices of the polytopes defined by the constrained matrix A . The geometric operation of moving from one vertex to another is called a pivot operation and the most commonly analyzed pivot operation with respect to smoothed analysis is the shadow vertex pivot method.

Without getting into technical details that will lead us too far afield, we note that the shadow vertex pivoting method requires us to calculate the following bound: let a_i for $1 \leq i \leq m$ denote the rows of the matrix A and let W be a fixed two dimensional subspace. We wish to bound

$$\mathbb{E}[|\text{edges}(\text{conv}(a_1, \dots, a_m) \cap W)|]$$

where $\text{conv}(a_1, \dots, a_m)$ is the convex hull of the rows (see [9] for more information).

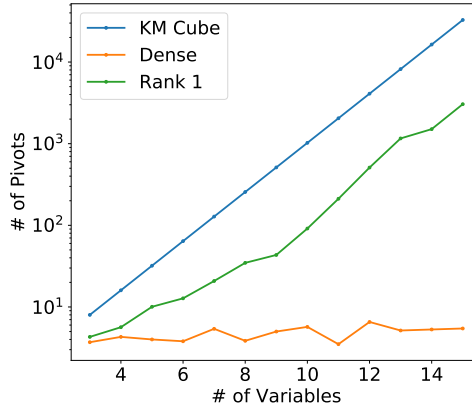
To calculate the above bound, we essentially need to understand the probability that $a_j^T \theta \leq t$ for a range of values of j and some $t \in \mathbb{R}$ (here θ represents the normal vector of the line connecting some two points a_i, a_k . For the pair a_i, a_k to be on the convex hull, we need the rest of the points to be on one side of the line). In the case that we add independent noise across the rows, this bound is possible to compute due to independence across a_j . However, in the case that we add rank 1 noise $u^T v$ (here $u \in \mathbb{R}^m, v \in \mathbb{R}^n$) to A , these probabilities become intractable using existing methods since a_j satisfying $a_j^T \theta \leq t$ gives us information about all other $a_{j'}$ for $j' \neq j$ since randomness is shared across the rows.

Nevertheless, it is possible to get a weak result for the smoothed analysis of the simplex method in our low-rank noise model by using a different pivoting operation. It is shown in [12] that if the rows satisfy a certain *geometric* property, then using a random pivoting rule results in an expected polynomial number of steps for the simplex method to converge.

The geometric property is the following: For any $I \subseteq [m]$, and $j \in [m]$, if a_j is not in the span generated by $a_i, i \in I$, then the distance from a_j to this span is at least δ . We note that the bound on the expected number of steps depends polynomially on $1/\delta$ and other parameters. This geometric property reduces to a singular value estimate as follows. For simplicity, lets focus on $j = 1$ and $I = \{2, \dots, n-1\}$. As in Section 1.3, it follows that $\|A_{[n]}^{-1} e_1\|$ is equal to $1/|w^T a_1|$ where w is the normal vector of the span of the rows a_2, \dots, a_n . Thus, if $s_n(A_{[n]})$ is “not too small” then $\|A_{[n]}^{-1} e_1\|$ cannot be “too large” and consequently, the distance from a_1 to the span of a_2, \dots, a_n is “not too small” (we are intentionally leaving our specific relations for a high level overview). The caveat is that we need the geometric property to hold between a_1 and *every* set of $n-1$ other vectors. However, since the bound of Theorem 13 only gives us an inverse polynomial probability, we cannot afford the union bound of $\binom{m}{n}$ unless $m = n + C$ for some constant C , which is not a realistic scenario.

Lastly, empirical evidence shows that rank 1 perturbation may not be a suitable if the *original* simplex method (the Dantzig simplex method) is used. In Figure 2, we use the Klee-Minty lower bound [14] for the Dantzig simplex method and add either a dense Gaussian or a rank 1 perturbation to the constraint matrix. We then plot the average number of pivot steps taken over twenty independent trials. It can be seen that a rank 1 perturbation only slightly improves over the exponential number of pivot steps required by the Dantzig simplex method whereas dense perturbations help greatly.

We conclude our discussion with a major open problem.



■ **Figure 2** The Dantzig simplex method applied to the Klee-Minty lower bound and its random perturbations.

► **Open Problem 26.** *Is there a pivoting rule for the simplex method that runs in expected polynomial time if we add random rank 1 noise to the constraint matrix?*

B.2 *k*-means clustering

Recall that in the *k*-means problem, we are given a set *X* of *n* points in \mathbb{R}^d and our goal is to partition the points into *k* sets S_i to minimize the objective

$$\sum_{i=1}^k \sum_{x \in S_i} \|x - \mu_i\|^2$$

where μ_i is the mean of the points in S_i . A common heuristic for this problem, also confusingly known as the *k*-means algorithm, or Lloyd’s method, is to randomly pick an initial set of *k* centers, assign each point in *X* to its closest center, update the means accordingly, and repeat until convergence. In the smoothed analysis framework, it was shown that if each point in *X* is perturbed by an independent Gaussian vector then convergence happens in polynomially many steps [3]. The existing analysis all crucially rely on the following geometric lemma.

► **Lemma 27.** *Let $x \in \mathbb{R}^d$ be drawn according to a *d*-dimensional Gaussian distribution of standard deviation σ , and let *B* be the *d*-dimensional ball of radius ε centered at the origin. Then $\mathbb{P}(x \in B) \leq (\varepsilon/\sigma)^d$.*

This lemma roughly states that the probability of a random Gaussian being in any ball of radius ε is at most ε^d , and is used to union bound over exponentially many events in the smoothed analysis of *k*-means.

Surprisingly, this lemma *does not* hold in our “rank 1” setting. More precisely, we can prove the following probabilistic bound which is a major impediment to understanding the smoothed complexity of the *k*-means problem with rank 1 noise.

► **Lemma 28.** *Let $x \in \mathbb{R}^d$ be drawn according to a standard *d*-dimensional Gaussian distribution and let $y \in \mathbb{R}$ be a scalar standard Gaussian random variable. If *B* is the *d*-dimensional ball of radius ε centered at the origin then $\mathbb{P}(yx \in B) = O(\varepsilon/\sqrt{d})$.*

Note that $yx \in \mathbb{R}^d$. We are considering random variables of this form because if a rank 1 perturbation was added to X , then each row is perturbed by a random vector of the form $yx \in \mathbb{R}^d$. Lemma 28 roughly states that the probability that the random vector yx is in any ball of radius ε only weakly depends on the dimension d . In particular, we do not get an exponentially small probability afforded by Lemma 27 that enables us to union bound over exponentially many events as in the arguments for the smoothed analysis of k -means under the standard noise model.

The intuition for Lemma 28 is as follows. First, note that from standard Gaussian concentration, we have $\|x\| \approx \sqrt{d}$. Treating this as fixed for now, this means that $y\|x\|$ is approximately distributed as a scalar Gaussian distribution with variance d . Therefore, from Proposition 12, it follows that $\mathbb{P}(\|y\|x\| \leq \varepsilon) = \Theta(\varepsilon/\sqrt{d})$. We now formalize this argument.

Proof. Note that $\|x\|_2^2$ is a chi-squared variable with d degrees of freedom. From [34], we know that the density of the product $Z = \|yx\|_2^2 = y^2\|x\|_2^2$ is given by

$$f_Z(z) \simeq \frac{1}{2^{d/2}\Gamma(d/2)} \int_0^\infty \left(x^{d/2-2} e^{-x/2}\right) \left(\frac{1}{\sqrt{z/x}} e^{-z/(2x)}\right) dx.$$

Therefore,

$$\mathbb{P}(\|yx\|_2^2 \leq \varepsilon^2) \simeq \frac{1}{2^{d/2}\Gamma(d/2)} \int_0^{\varepsilon^2} \int_0^\infty \left(x^{d/2-2} e^{-x/2}\right) \left(\frac{1}{\sqrt{z/x}} e^{-z/(2x)}\right) dx dz.$$

We now switch the order of summation which is valid since the integrand is positive. From the definition of the error function, we can check that

$$\int_0^{\varepsilon^2} \frac{1}{\sqrt{z/x}} e^{-z/(2x)} dz \simeq x \cdot \operatorname{erf}(\varepsilon/\sqrt{x}).$$

We now use the estimate $\operatorname{erf}(t) \leq 2t$ which holds for all $t \geq 0$. This gives us

$$\int_0^{\varepsilon^2} x^{d/2-1} e^{-x/2} \operatorname{erf}(\varepsilon/\sqrt{x}) dx \lesssim \varepsilon \int_0^{\varepsilon^2} x^{d/2-3/2} e^{-x/2} dx = \varepsilon 2^{d/2-1/2} \Gamma(d/2 - 1/2).$$

Finally, noting that $\Gamma(d/2 - 1/2)/\Gamma(d/2) \lesssim 1/\sqrt{d}$ gives us our desired probability bound. \blacktriangleleft

Note that the above bound is the best that we can hope for. Indeed, we can say that $\|x\|_2^2 = \Omega(d)$ with probability $1/2$ so conditioning on this event, we have that $\Pr(\|y\|x\|_2 \leq \varepsilon) = \Omega(\varepsilon/\sqrt{d})$. We note that Lemma 27 is also required for the smoothed analysis of other Euclidean problems such as a local search heuristic for Euclidean TSP [17].