# Smoothed analysis of the condition number under low-rank perturbations

Rikhav Shah* Sandeep Silwal†

**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. We experimentally validate our approach and consider generalizations to symmetric and complex settings. Lastly, we show barriers in applying our low rank model to other problems studied in the smoothed analysis framework.

## 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 worse-case and average case analysis [14]. In this framework, we are given an arbitrary input that is then perturbed randomly according to some 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 behaviour of many numerical algorithms.

To give context to our result, recall that the condition number of a $n \times n$ matrix $M$ with singular values $s_1(M) \geq \cdots \geq s_n(M)$ is defined as the ratio $s_1(M)/s_n(M)$. Generally, we say that 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 2.1 if the entries are not too large compared to the size of $M$. Therefore, the bulk of the work lies in controlling the smallest singular value $s_n(M)$. Extending a result of Edelman [5], Sankar, Speilman and Teng showed the following result in [11]:

**Theorem 1.1.** *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 $t > 0$,*

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

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 [19, 15].

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 noise and less space. Our results demonstrate the answer is yes.

*Email: rikhav.shah@berkeley.edu.
†Email: silwal@mit.edu.

## 1.1 Our results

Our primary technique is to replace the Gaussian ensemble $N_n$ with a *low rank* matrix.

**Theorem 1.2** (Theorem 3.1 simplified)**.** *Let $1 \le k \le n/2$ 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*

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

*for absolute constants $C, c > 0$.*

**Theorem 1.3** (Theorem 3.3 simplified)**.** *Let $1 \le k \le n/2$ 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*

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

*for absolute constants $C, c > 0$.*

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 1.2 in Section 3 and discuss the dependence of the term $s_{n-k}(M)$ which we show is unavoidable.

Theorem 1.2 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. Later in Section 3.1 we show that our result cannot hold if we pick the entries of $U, V$ to be from general sub-Gaussian distributions. Additionally in Section 4, we present numerical evidence for our low rank error model. Finally in Section 5, 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. We show that current analysis methods that work for dense random perturbations for these problems do not carry over to the low rank case.

**Remark 1.4.** We note that Theorem 1.2 requires the input matrix $M$ to have rank exactly $n-k$. This is just a technicality to achieve clean bounds and can be circumvented with the use of Weyl's perturbation inequality which implies that if $M + E = M'$, then $|s_n(M) - s_n(M')| \le \|E\|$. For example in the case of $k = 1$, if $M$ is of full rank but contains one small singular value $\sigma_n(M)$, then we can let $E = -\sigma_n(M) u_n v_n^T$ where $u_n, v_n$ are the left and right singular vectors corresponding to $\sigma_n(M)$. We can then apply Theorem 1.2 to $M'$ and arrive at a bound for the smallest singular value of the perturbation of $M$ up to error $O(\sigma_n(M))$. Since our ideal use case is when $\sigma_n(M)$ is already negligible, the final bound that we get is comparable to the bound from Theorem 1.2. We show that this technicality does not impact experimental performance in Section 4.

## 1.2 Previous techniques and our approach

In summary, previous techniques do not work in our case since we have shared independence across different rows/columns of the matrix. In more detail, all of the previous techniques used to bound the singular values of a deterministic matrix plus a random noise (or just 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.

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, since 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 inequality and its variants.

To be more precise, lets consider a high level overview of the proof of Theorem 1.1. 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 $r_2, \cdots, r_n$. Then the proof proceeds by noting that $w^T r_1 = \mathcal{N}(z, 1)$ for some $z \in \mathbb{R}$ and then standard anti-concentration results of $\mathcal{N}(0, 1)$ can be used. In the more general case of Tao and Vu [19, 15], 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 simply cannot fix a normal hyperplane of a span of a subset of rows since this automatically implies something about the rows not considered in the span. Therefore, most of the spectrum of existing techniques are not applicable in our case.

To overcome these barriers, we first reduce our problem to adding noise to a diagonal matrix. 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.

## 1.3 Mathematical motivation for low rank model

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 $s_1(D), \cdots, s_{n-1}(D)$ while the last entry is 0. Then, the adjugate matrix is is the all zeros matrix except the top leftmost entry which is $s_1(D) \cdots s_{n-1}(D)$. Therefore,

$$\det(D) + v^T \text{adj}(D) u = (u_1 v_1)(s_1(D) \cdots s_{n-1}(D))$$

which is non-zero with probability 1 since $u_1 v_1 \neq 0$ with probability 1. Thus, adding a random rank 1 perturbation results in $D$ not being singular which motives the question of studying the smallest singular value after a random rank 1 (and more generally low rank) perturbation.

## 1.4 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, 1], the perceptron algorithm [3], and the simplex method [14, 4]. In all of these results, the goal is to show that after an input instance of the problem is suitable 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 [17, 13, 8] 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 [11] 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.

### 1.5 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 $C_1, C_2$ denote positive constants.

## 2 Preliminary results

In this section we enumerate some useful results. First, we recall a classical estimate of the operator norm of a random matrix of Seginer [12]. The following proposition essentially shows that the top singular value of a random matrix is well behaved under mild assumptions.

**Proposition 2.1.** *Let $M$ be a random $n \times n$ matrix with entries $m_ij$. Then,*

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

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

**Lemma 2.2** (Theorem 1.1 reformulated.)**.** *Let $G \in \mathbb{R}^{k\times k}$ with all entries chosen i.i.d. from $\mathcal{N}(0,1)$. Then*

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

*for some absolute constant $C$.*

**Lemma 2.3** (Theorem 1.1 in [16])**.** *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*

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

**Lemma 2.4** (Proposition 2.3 in [10])**.** *Let $G \in \mathbb{R}^{(n-k)\times k}$ for $k \le n/2$ with all entries chosen i.i.d. from $\mathcal{N}(0,1)$. Then*

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

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

**Lemma 2.5.** *Let $G \in \mathbb{C}^{(n-k)\times k}$ for $k \le n/2$ with all entries chosen with i.i.d. real and imaginary parts from $\mathcal{N}(0,1/2)$. Then*

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

*for $t_2, C_1, C_2$ as in Lemma 2.4.*

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

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

where the last inequality follows by Lemma 2.4 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 2.6.** *Let*

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

*be an $n \times n$ matrix. Then*

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

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

*Proof.* We first use the 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.

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

□

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

**Proposition 2.7.** *Let $x \sim \mathcal{N}(0,1)$. Then, $\mathbb{P}(|x| \leq \varepsilon) = O(\varepsilon)$.*

# 3 Proof of main theorems

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

**Theorem 3.1.** *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*

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

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

Before we present the proof, let us briefly mention why the term $s_{n-k}(M)$ is unavoidable. For simplicity, consider $k = 1$ and suppose that $M$ is of rank $n-1$ and suppose its smallest nonzero singular value is equal to $\delta$. After adding a rank 1 term $uv^T$ to $M$, its rank is $n$ with probability 1. However, if we consider the limit $\delta \to 0$, then $M + uv^T$ approaches a rank $n-1$ matrix meaning $s_n(M + uv^T) \to 0$. Hence, any concentration bound such as (1) must depend on the term $s_{n-1}$. We now proceed to the proof.

Our strategy to prove Theorem 3.1 will reduce general $M$ to the case of $M$ nonnegative and diagonal, then express $s_n(M + UV^T) = s_1((M + UV^T)^{-1})$ 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.

**Lemma 3.2.** *Let $D = \mathrm{diag}(s_n(M), \cdots, s_1(M))$. Let $U, V$ be as in Theorem 3.1. 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 3.1.* 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 3.2 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 2.6 to upper bound $s_n(M + UV^T)$. The factor corresponding to the Schur complement is

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

and the resulting bound is

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

Denote events

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

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

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

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

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

Taking the reciprocal of both sides yields

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

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:

$$\begin{aligned} \mathbb{P}(\neg\mathcal{E}_1 \vee \neg\mathcal{E}_2) &\le \mathbb{P}(\neg\mathcal{E}_1) + \mathbb{P}(\neg\mathcal{E}_2) \\ &\le \mathbb{P}(s_1(U_{[n]\setminus[k]}) \ge t_2\sqrt{n-k}) + \mathbb{P}(s_1(V_{[n]\setminus[k]}) \ge t_2\sqrt{n-k}) \\ &\quad + \mathbb{P}(s_k(U_{[k]}) \le t_1/\sqrt{k}) + \mathbb{P}(s_k(V_{[k]}) \le t_1/\sqrt{k}) \\ &\le 2\,C_1 t_1 + 2\,C_2 e^{-C_3\,t_2^2\,n}. \end{aligned}$$

where the last step follows by applying Lemmas 2.2 and 2.4 twice each. The factors of 2 can be subsumed into the constants $C_1$ and $C_2$ giving the final result. □

**Theorem 3.3.** *Let $M, t_1, t_2, C_2, C_3$ be as in theorem 3.1. Let $U, V$ be $n \times k$ complex matrices with real and imaginary parts drawn independently from $\mathcal{N}(0, 1/2)$. Then*

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

Note that the first term on the righthand side is $2t_1^2$ rather than $C_1 t_1$ as it was in Theorem 3.1.

*Proof.* The only place the proof differs from the proof of Theorem 3.1 is in the upper bound on $\mathbb{P}(\neg \mathcal{E}_1)$. Instead of $C_1 t_1$, it is simply $t_1^2$ by Lemma 2.3. □

**Remark 3.4.** Theorems 3.1 and 3.3 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 3.2, 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 $\mathbb{P}(\neg \mathcal{E}_1 \vee \neg \mathcal{E}_2)$ by a factor of 2. □

## 3.1 Sub-Gaussian perturbations

It is of interest to generalize Theorem 3.1 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 $\pm 1$ (Rademacher) random variables. We show in this case that we cannot state a general statement unless extra assumptions about the fixed matrix $M$ is made.

For simplicity, consider the case where we add a rank 1 random matrix $uv^T$ to $M$ where the entries of $u, v \in \mathbb{R}^n$ are independent Rademacher random variables and $M$ has rank $n-1$. Our goal is to give a lower bound on $s_n(M + uv^T)$. As in the proof of Theorem 3.1, 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 $\pm 1$. Here, the properties of $L, R$ can have substantial impact on $s_n(M + uv^T)$. For example, we can suppose that the top left entry of $D$ is 0. Then, if the first row of $L$ is *sparse*, 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 0 which implies that $M + uv^T$ is still rank $n-1$ with constant probability. Therefore, a general statement such as Theorem 3.1 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 3.1 carries through in the $\pm 1$ 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.

## 3.2 Applications 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 [18].

**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 decent method. If the conjugate gradient decent 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 decent method are required.

The usefulness of our low rank error model is further supported by the conjugate gradient decent 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 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 [7] 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 $C \approx 1.47$. This implies that the smallest singular value of $M_n$ is also at most $O(n/C^n)$.

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 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.

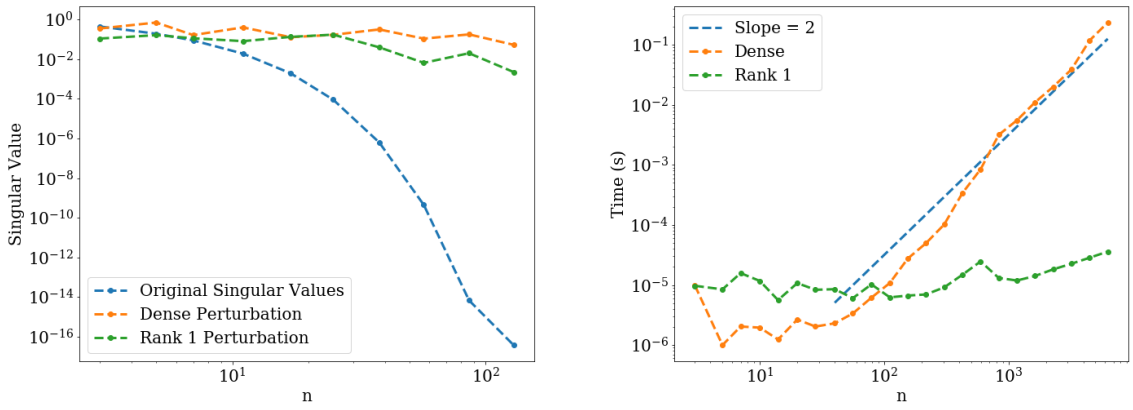


(a) (b)

Figure 1: (a) Smallest singular values of the original matrix and dense and rank 1 perturbations. (b) Time taken to perform a matrix vector product after a dense perturbation and a rank 1 perturbation.

# 5 Challenges for other Problems

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.

## 5.1 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

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

where $\text{conv}(a_1, \cdots, a_m)$ is the convex hull of the rows (see [4] 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 $\theta$ 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 [6] 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 $\delta$. 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 $I = \{2, \cdots, n-1\}$. As in Section 1.2, 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 $a_2, \cdots, a_n$. 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 $a_2, \cdots, a_n$ 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 3.1 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. We conclude our discussion with a major open problem.

**Open Problem 5.1.** 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?

## 5.2 $k$-means

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 $\mu_i$ 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 [1]. The existing analysis all rely crucially on the following geometric lemma.

**Lemma 5.2.** *Let $x \in \mathbb{R}^d$ be drawn according to a $d$-dimensional Gaussian distribution of standard deviation $\sigma$, and let $B$ be the $d$-dimensional ball of radius $\varepsilon$ centered at the origin. Then $\mathbb{P}(x \in B) \leq (\varepsilon/\sigma)^d$.*

Surprisingly, the above 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 5.3.** *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 $\varepsilon$ 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 5.3 roughly states that the probability that the random vector $yx$ is in any ball of radius $\varepsilon$ only weakly depends on the dimension $d$. In particular, we do not an exponentially small probability afforded by Lemma 5.2 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 5.3 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 2.7, 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 [20], 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

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

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

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

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

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 5.2 is also required for the smoothed analysis of other Euclidean problems such as a local search heuristic for Euclidean TSP [9].

# References


[1] D. Arthur, B. Manthey, and H. Röglin. Smoothed analysis of the k-means method. *J. ACM*, 58(5), Oct. 2011.

[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] A. Blum and J. 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.

[4] D. Dadush and S. 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.

[5] A. Edelman. Eigenvalues and condition numbers of random matrices. *SIAM Journal on Matrix Analysis and Applications*, 9(4):543–560, 1988.

[6] F. Eisenbrand and S. S. Vempala. Geometric random edge. *Math. Program.*, 164(1-2):325–339, 2017.

[7] N. D. E. (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).

[8] B. Manthey and H. Raglin. Smoothed analysis: Analysis of algorithms beyond worst case. *it - Information Technology*, 53(6):280 – 286, 2011.

[9] B. Manthey and R. Veenstra. Smoothed analysis of the 2-opt heuristic for the tsp: Polynomial bounds for gaussian noise. In L. Cai, S.-W. Cheng, and T.-W. Lam, editors, *Algorithms and Computation*, pages 579–589, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[10] M. Rudelson and R. Vershynin. Smallest singular value of a random rectangular matrix. *Communications on Pure and Applied Mathematics*, 62(12):1707–1739, 2009.

[11] A. Sankar, D. A. Spielman, and S.-H. Teng. Smoothed analysis of the condition numbers and growth factors of matrices. *SIAM Journal on Matrix Analysis and Applications*, 28(2):446–476, 2006.

[12] Y. Seginer. The expected norm of random matrices. *Combinatorics, Probability and Computing*, 9:149–166, 03 2000.

[13] D. A. Spielman and S. hua Teng. Smoothed analysis: an attempt to explain the behavior of algorithms in practice. *COMMUN. ACM*, 2009.

[14] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *J. ACM*, 51(3):385–463, May 2004.

[15] T. TAO and V. VU. Smooth analysis of the condition number and the least singular value. *Mathematics of Computation*, 79(272):2333–2352, 2010.

[16] V. Tau. Random matrices: the distribution of the smallest singular values. *Geometric and Functional Analysis*, 20:260–297, 03 2010.

[17] S.-H. Teng. Smoothed analysis of algorithms and heuristics. In L. Wang, editor, *Computing and Combinatorics*, pages 10–11, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

[18] L. N. Trefethen and D. I. Bau. *Numerical linear algebra*. SIAM Society for Industrial and Applied Mathematics, 2000.

[19] V. H. Vu and T. 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.

[20] Wikipedia contributors. Product distribution — Wikipedia, the free encyclopedia, 2020. [Online; accessed 18-August-2020].