

Matrix Perturbation: Davis-Kahan in the Infinity Norm

Abhinav Bhardwaj*

Van Vu†

Abstract

Perturbation theory is developed to analyze the impact of noise on data and has been an essential part of numerical analysis. Recently, it has played an important role in designing and analyzing matrix algorithms. One of the most useful tools in this subject, the Davis-Kahan sine theorem, provides an ℓ_2 error bound on the perturbation of the leading singular vectors (and spaces).

We focus on the case when the signal matrix has low rank and the perturbation is random, which occurs often in practice. In an earlier paper, O’Rourke, Wang, and the second author showed that in this case, the Davis-Kahan bound can be improved significantly. In particular, the noise-to-gap ratio condition in the original setting can be weakened considerably.

In the current paper, we develop an infinity norm version of the O’Rourke-Vu-Wang result. The key ideas in the proof are a new bootstrapping argument and the so-called iterative leave-one-out method, which may be of independent interest.

Applying the new bounds, we develop new and very simple algorithms for several well-known problems, such as finding hidden partitions and matrix completion. The core of these new algorithms is the fact that one is now able to quickly approximate certain key objects in the infinity norm, which has critical advantages over approximations in the ℓ_2 norm, Frobenius norm, or spectral norm.

1 Introduction

1.1 The classical Davis-Kahan theorem. Perturbation theory is developed to analyze the impact of noise on data and has been an essential part of numerical analysis. The general setting is that we have a (signal or data) matrix A , a noise matrix E , and a matrix functional f . Our goal is to compare $f(A)$ with $f(A + E)$. A typical perturbation bound provides an upper bound for the difference $f(A + E) - f(A)$ in some norm.

For the sake of presentation, in most of the paper, we assume that both A and E are symmetric and of dimension n . All results in this paper can be extended to the asymmetric case by a simple symmetrization trick. We assume that n is sufficiently large, whenever needed, and asymptotic notations are used under the assumption that n tends to infinity.

Assume that A has rank r and let σ_i be the non-trivial singular values of A in decreasing order, for $1 \leq i \leq r$. Let $\mathbf{u}_i$ be the corresponding singular vector of σ_i with entries u_{il} . For the sake of presentation, we assume that σ_i are different so $\mathbf{u}_i$ are well defined, up to sign. Let $\tilde{A} = A + E$, and use notation $\tilde{\sigma}_i$, $\tilde{\mathbf{u}}_i$, and $\tilde{u}_{il}$. Notice that because $\mathbf{u}_i$ and $\tilde{\mathbf{u}}_i$ are unique up to sign, we may always choose the signs so that the angle between them is at most $\pi/2$. Let $\Delta_i = \sigma_i - \sigma_{i+1}$, and let $\delta_i = \min\{\Delta_{i-1}, \Delta_i\}$ be the distance from σ_i to the nearest singular value. We take $\sigma_0 = \infty$ for the sake of consistency.

One of the most useful tools in perturbation theory is the Davis-Kahan bound, which provides a perturbation bound for singular vectors.

THEOREM 1.1. (DAVIS-KAHAN) *There is a constant $C > 0$ such that, provided $\delta_i \geq 2\|E\|$,*

$$(1.1) \quad \|\mathbf{u}_i - \tilde{\mathbf{u}}_i\|_2 \leq C \frac{\|E\|}{\delta_i}.$$

The first version of this theorem, by Davis and Kahan [29] was stated for eigenvectors (and eigenspaces). Later, Wedin [62] extended the results to singular vectors. In this paper, we use singular vectors, which are simpler to handle. The more general version of the Davis-Kahan theorem gives a perturbation bound for the spaces spanned by a set of singular vectors. In this paper, we focus on individual singular vectors, but the results can be extended into that direction with simple modifications.

*Department of Mathematics, Yale University.

†Department of Mathematics, Yale University.

It is important to notice that for the RHS of (1.1) to be small, one needs

$$(1.2) \quad \gamma \|E\| \leq \delta_i,$$

for some large $\gamma > 0$. In other words, the noise to gap ratio $\frac{\|E\|}{\delta_i}$ has to be small. We will refer to this as the *noise-to-gap ratio* assumption.

Notation. We use the conventional asymptotic notations, such as o, O, Ω, Θ . We will also use the notation $f(n) = \tilde{O}(g(n))$ if there exists some absolute c such that $f(n) = O(g(n) \log^c n)$; $\tilde{\Theta}$ and $\tilde{\Omega}$ are defined similarly.

1.2 Low rank data with random noise and an improved version of Davis-Kahan theorem. In modern studies, the following two assumptions come up frequently. First the data matrix A has low rank, and second, the noise matrix E is random.

The low rank (or approximately low rank) phenomenon is automatic in a number of theoretical settings, such as the clustering problem discussed in Section 4. It also occurs in so many real life problems that researchers have even tried to give a theoretical explanation for this; see [57].

Under the assumption that A has low rank r and E is random, the second author discovered that one can improve the Davis-Kahan bound [60]. In particular, one can replace the noise-to-gap ratio assumption (1.2) by much weaker ones; for related results, see [1, 3, 9, 11, 23, 36, 46, 49, 51, 52, 60, 64].

ASSUMPTION 1.1. *In what follows, we assume that A is a symmetric, deterministic matrix with rank r . E will be a random symmetric matrix with independent (but not necessarily iid) upper triangular entries ξ_{ij} . The ξ_{ij} will be K -bounded random variables with mean 0. A random variable ξ is K -bounded if $|\xi| \leq K$ with probability 1.*

Following [60], about 10 years ago, O’Rourke, Wang, and the second author [51] obtained the following theorem.

THEOREM 1.2. *For any constants $\tau, r > 0$, there exists a constant C_0 such that with probability at least $1 - \tau$,*

$$(1.3) \quad \|\tilde{\mathbf{u}}_1 - \mathbf{u}_1\|_2 \leq C_0 \left[\frac{Kr^{1/2}}{\delta_1} + \frac{\|E\|}{\sigma_1} + \frac{\|E\|^2}{\delta_1 \sigma_1} \right].$$

The theorem holds for other singular vectors as well, with simple modifications. A more quantitative form of this theorem [51] allows one to take $\tau \rightarrow 0$ with n .

The key point in Theorem 1.2 is that for the RHS to be small, we only need to assume $\frac{\|E\|^2}{\sigma_i \delta_1}$ and $\frac{1}{\delta_1}$ are small, which is much weaker than (1.2), where we need to require that $\frac{\|E\|}{\delta_1}$ is small. (In standard settings, $\|E\|$ is of order $\sqrt{n}$).

In many applications (see [51]), the gap δ_1 is smaller than $\|E\|$, so the “noise-to-gap ratio is small” assumption (1.2) is violated. On the other hand, even if $\delta_1 \leq \|E\|$, it is still often the case that the product $\delta_1 \sigma_1$ is larger than $\|E\|^2$, as σ_1 can be way larger than both δ_1 and $\|E\|$, and our bound applies. Let us illustrate this with an example.

Example. Let A be a matrix whose entries are of order $\Theta(1)$ with constant rank r , and E be a matrix whose entries are iid standard Gaussian. Since $\sum_{i=1}^r \sigma_i^2 = \|A\|_F^2 = \Theta(n^2)$, we expect that the non-trivial singular values of A are of order $\Theta(n)$. On the other hand, it is well known that (with high probability), $\|E\| = (2 + o(1))\sqrt{n}$. Furthermore, by a simple truncation trick, we can set $K = 20\sqrt{\log n}$, as this holds with overwhelming probability. Thus, in Theorem 1.2, we only need to require the gap δ_i to be $\Omega(\sqrt{\log n})$ to have a meaningful conclusion (making the RHS of (1.3) going to zero). On the other hand, an application of Davis-Kahan theorem would require $\delta_i = \Omega(\sqrt{n})$, a significantly stronger assumption, to achieve the same conclusion.

For more recent progress in this direction, see [52].

2 The infinity norm version

Theorem 1.1 provides an ℓ_2 estimate. It is natural and important to obtain similar results in the infinity norm. Going from ℓ_2 to ℓ_∞ is always a non-trivial task and progress has only been made in the last 10 years or so.

There are many papers in this topic considering either the infinity norm or the $\ell_{2 \rightarrow \infty}$ norm [4, 5, 17, 22, 36, 37]. However, in all of these papers, one needs to use the original noise to gap ratio assumption (1.2). Our study in this paper will go well beyond this setting, as our goal is to obtain an infinity norm version of Theorem 1.2, which holds under weaker assumptions.

While finishing this paper, we became aware of a result in [27]. In this paper, the authors studied a *hybrid* model where a *symmetric* low rank matrix is perturbed with *asymmetric* random noise. This paper also studied the infinity norm and does not need to assume (1.2). On the other hand, the analysis relies strongly on the hybrid model and is totally different from the methods in this paper; see [27] for details. The hybrid model does not seem to occur very often in applications, as it is natural to assume that A and E have the same type of symmetry.

3 New results

3.1 An optimal guess. Our goal is to find an infinity norm analogue of Theorem 1.2, the improved version of Davis-Kahan theorem, with a significantly weakened noise to gap assumption. Consider a singular vector $\mathbf{u}_i$, its perturbed counterpart $\tilde{\mathbf{u}}_i$, and the infinity norm difference $\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty$. To start, let us raise a question.

QUESTION 3.1. *What is the best possible bound we can hope to achieve for $\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty$?*

Consider the ℓ_2 norm difference $\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2$. It is apparent that

$$\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty \geq \frac{1}{\sqrt{n}} \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2.$$

Thus, the best bound one would hope for (up to a poly-logarithmic factor, which is usually unavoidable in a random setting) is

$$(3.4) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty = \tilde{O}\left(\frac{1}{\sqrt{n}} \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2\right).$$

However, (3.4) may be too optimistic. In practice, it is natural to expect that coordinate-wise errors are proportional to the magnitude of the coordinates. Simply speaking, the error at a higher magnitude coordinate is likely to be larger. Thus, a more realistic version of (3.4) is

$$(3.5) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty = \tilde{O}(\|\mathbf{u}_i\|_\infty \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2).$$

We are going to prove that under certain conditions, a slightly weaker version of (3.5) holds. The precise form of the result is a bit technical, but in essence it shows (see Remark 3.3 for discussion)

$$(3.6) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty = \tilde{O}(\|U\|_\infty \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2).$$

The parameter $\|U\|_\infty$ is not an adhoc one. It has played an important role in many applications of spectral methods in statistics, through the notion of *incoherence* [20, 21, 25, 54]. Our main theorem roughly asserts that under a modest condition on the singular values and the gaps, (3.6) holds. Let us illustrate with a corollary of our main results.

THEOREM 3.1. (LEADING SINGULAR VECTOR PERTURBATION) *Let E be a symmetric, K -bounded random matrix with independent upper triangular entries. Then with probability $1 - o(1)$,*

$$(3.7) \quad \|\tilde{\mathbf{u}}_1 - \mathbf{u}_1\|_\infty \leq c \|U\|_\infty \left[\frac{\|E\|}{\sigma_1} + \frac{K\sqrt{\log n}}{\delta_1} + \frac{\|E\|^2}{\sigma_1 \delta_1} \right] + \frac{cK\sqrt{\log n}}{\sigma_1}.$$

Notice that in the main term $c \|U\|_\infty \left[\frac{\|E\|}{\sigma_1} + \frac{K\sqrt{\log n}}{\delta_1} + \frac{\|E\|^2}{\sigma_1 \delta_1} \right]$, the term $\left[\frac{\|E\|}{\sigma_1} + \frac{K\sqrt{\log n}}{\delta_1} + \frac{\|E\|^2}{\sigma_1 \delta_1} \right]$ is essentially the RHS of the bound (1.3) for $\|\tilde{\mathbf{u}}_1 - \mathbf{u}_1\|_2$. In the case K grows slowly with n , the second term $\frac{cK\sqrt{\log n}}{\sigma_1}$ is often negligible compared to the main term.

In the next three sections, we present our main theorems.

3.2 Main theorems: The deterministic setting. Our first main theorem is a deterministic one (Theorem 3.2) where we can measure the difference of the eigenvectors between two matrices A and $A + H$, where both A and H are deterministic. This theorem asserts a relation between the eigenvectors of A and $A + H$ through information on the eigenvalues of A , $A + H$ and those of the principal minors of $A + H$. This is somewhat close, in spirit, to the eigenvector-eigenvalue identity, discovered several times in linear algebra, most recently through the study of neutrino oscillations by Denton, Parke, Tao, and Zhang [32]; see [31] for a survey.

We keep the definition of all parameters such as $\epsilon_1(i), \epsilon_2(i)$ (with H playing the role of E). We denote by $H^{\{l\}}$ the matrix obtained by zeroing out the l th row and column of H .

Let A and H be symmetric matrices of size n , where A has rank r . For any $1 \leq l \leq n$, set $A^{\{l\}} = A + H^{\{l\}}$ and $\tilde{A} = A + H$. Notations such as $\sigma_i, \tilde{\sigma}_i$ and $\sigma_i^{\{l\}}$ are self-explanatory.

THEOREM 3.2. *Consider $A, H, \tilde{A}, A^{\{l\}}, H^{\{l\}}$ as above. Let $U^{\{l\}}$ denote the $n \times r$ matrix of r leading singular vectors of $A^{\{l\}}$ and $\mathbf{x} = \mathbf{x}(l)$ be the l th row of H , except with the l th entry of $\mathbf{x}$ reduced to $H_{ll}/2$. Set $C_0 := 272 \times 4r^{3/2}$, and define $a_l := \|U^{\{l\}T} \mathbf{x}\|_2$. Suppose that*

- $\sigma_i > C_0 \|H\|$
- $\delta_i > C_0 \max\{a_l, \kappa_i \|H\| \|U\|_\infty\}$
- $\min\{|\tilde{\sigma}_i - \sigma_{i+1}^{\{l\}}|, |\tilde{\sigma}_i - \sigma_{i-1}^{\{l\}}|\} > \delta_i/2$.

Then,

$$(3.8) \quad |\tilde{u}_{il} - u_{il}| \leq C_0 \|U_{l,\cdot}\|_\infty \left[\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + a_l \kappa_i \epsilon_2(i) \right] + 256r \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i},$$

where $\tilde{\mathbf{u}}_i$ is the i th singular vector of $A + H$.

REMARK 3.1. *The first assumption $\sigma_i > C_0 \|H\|$ is a signal to noise assumption.*

The second assumption $\delta_i > C_0 \kappa_i r^{1/2} \max\{a_l, \|H\| \|U\|_\infty\}$ is a gap assumption (replacing the stronger assumption (1.2) from the original Davis-Kahan theorem). In many applications (including all applications in this paper), $\kappa_i = O(1)$ and $\|U\|_\infty = n^{-1/2+o(1)}$, thus $\kappa_i \|H\| \|U\|_\infty = \|H\| n^{-1/2+o(1)}$, improving (1.2) by a factor of $n^{-1/2+o(1)}$. The term a_l measures the correlation between H and A . This is small if $\mathbf{x}$ does not align with any non-trivial eigenvector of $A^{\{l\}}$. If H is random then $\mathbf{x}$ is a random vector independent of A and this holds trivially.

The last assumption is a stability assumption. Intuitively, we expect that $\tilde{\sigma}_i$ is close to σ_i and $\sigma_{i+1}^{\{l\}}$ close to σ_{i+1} , which would imply that $\tilde{\sigma}_i - \sigma_{i+1}^{\{l\}}$ is close to $\sigma_i - \sigma_{i+1} \geq \delta_i$. Our stability assumption guarantees a weaker bound $|\tilde{\sigma}_i - \sigma_{i+1}^{\{l\}}| \geq \delta_i/2$.

3.3 Main theorems: The random setting with small K . Now we consider the random model $A + E$, where E is random matrix whose entries are K bounded random variable. The result in this section holds for any K , but for large K , we have a better result (under a slightly stronger assumption) in the next section.

To ease the presentation, we introduce the following definition.

DEFINITION 3.1. *For a matrix A , we say that a singular value and its gap (σ_i, δ_i) is (c, τ, ν) stable under E if the following are all true. Let $T = \inf\{t > 0 : \mathbb{P}(\|E\| > t) \leq \tau\}$.*

- (a) $\sigma_i > cT$.
- (b) $\delta_i > c(K \log^{\nu/2} n + \sigma_i^{-1} T^2)$.
- (c) $\delta_i > c \kappa_i T \|U\|_\infty$; where $\kappa_i := \sigma_1 / \sigma_i$.

The conditions in this definition will guarantee that σ_i, δ_i are stable, in that they do not vary too much after the perturbation by E . This gives us control on $\tilde{\sigma}_i, \tilde{\delta}_i$, which is important in the analysis. Most importantly, it guarantees that the stability assumption in Theorem 3.2 hold; see Remark 3.1.

We will assume that our signal matrix A has a singular value and gap (σ_i, δ_i) that is (c, τ, ν) stable under E for properly chosen c, τ, ν , where τ is a parameter that goes to zero, ν is a small constant (like 1 or 2), and c is a large constant.

Let us comment on each condition. We can think of T as basically $\|E\|$, since T is a stand-in for a high probability bound of $\|E\|$, which is often strongly concentrated; see [59].

REMARK 3.2. (INTERPRETATION OF STABILITY) *The conditions here run parallel with those in Theorem 3.2.*

- The first condition is simply the assumption that the signal-to-noise ratio is large. This is absolutely necessary because if the intensity of the noise is larger than that of the signal, the signal will most likely be destroyed [53].
- The second condition essentially asks for the gap to be at least polylogarithmic and the product $\delta_i \sigma_{i+1}$ to dominate $\|E\|^2$ (which is consistent with the improved ℓ_2 bound in Theorem 1.2).
- The third condition requires the gap to be at least $\kappa_i \|U\|_\infty \|E\|$. This is better than (1.2) by a factor $\kappa_i \|U\|_\infty$, which can be as small as $n^{-1/2}$, a large improvement. As a matter of fact, in all applications in this paper, $\kappa_i = O(1)$ and $\|U\|_\infty = n^{-1/2+o(1)}$.

As the role of E is consistent through the paper, we will simply say (σ_i, δ_i) is (c, τ, ν) stable, instead of saying that (σ_i, δ_i) is (c, τ, ν) stable under E .

Set $\epsilon_1(i) := \|E\|/\sigma_i$, $\epsilon_2(i) := 1/\delta_i$, $C(r) = 1000 \times 9^{2r}$, $c_1(r) = 2500r^{3/2}$. (Big constants like 1000 are for definiteness and are rather adhoc; we make no attempt to optimize these in the current paper.) We are going to use these parameters throughout the paper.

THEOREM 3.3. *Let $c_0, \tau > 0$, where c_0 is a constant and τ may tend to zero with n . Set $c = 2^{11}(c_0 + 1)r^3$. Assume that (σ_i, δ_i) is $(c, \tau, 1)$ stable. Then with probability at least $1 - C(r)n^{-c_0} - 2\tau$,*

$$(3.9) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty \leq c \|U\|_\infty (\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + \kappa_i \epsilon_2(i) K \sqrt{\log n}) + \frac{cK \sqrt{\log n}}{\sigma_i}.$$

Theorem 3.1 follows easily from Theorem 3.3 and Theorem 1.2

REMARK 3.3. (OPTIMALITY) *Consider the bound in Theorem 3.3 for the first singular vector. Let us compare it to the desired bound (3.6), which is $\|\tilde{\mathbf{u}}_1 - \mathbf{u}_1\|_\infty = \tilde{O}(\|U\|_\infty \|\tilde{\mathbf{u}}_1 - \mathbf{u}_1\|_2)$. We notice that the bound (3.9) is off by 2 terms: $\frac{cK \sqrt{\log n}}{\sigma_i}$ and $\tilde{O}(\|U\|_\infty (\epsilon_1 + \epsilon_2))$.*

In many applications, σ_1 is sufficiently large and K is sufficiently small that the first term $\frac{cK \sqrt{\log n}}{\sigma_1}$ is negligible. Furthermore, when applying (3.9), we typically do not know $\|\tilde{\mathbf{u}}_1 - \mathbf{u}_1\|_2$. In this case, the best we can do is to use the bound from Theorem 1.2, which contains both ϵ_1 and ϵ_2 . Moreover, a new study [52] reveals that both ϵ_1 and ϵ_2 are necessary in (1.3). Thus, technically (3.9) has achieved what (3.6) promised. The case when K is relatively large will be discussed in the next section.

Our proof also reveals that if we consider the local estimate $|\tilde{u}_{il} - u_{il}|$, then we can replace $\|U\|_\infty$ by $\|U_{l,\cdot}\|_\infty$ where $U_{l,\cdot}$ is the l th row of U . So for a local estimate, we only need local information from U .

3.4 Main theorems: Random setting with large K . Let us discuss the last term $K \sqrt{\log n}/\sigma_i$ in (3.9). We stated that in many cases, this term is negligible. It is indeed so when the entries of E have a fixed distribution, which does not depend on n , as illustrated by the following two examples.

Example. If the entries of E are iid Rademacher (± 1) , then $K = 1$.

Example. If the entries of E are iid $N(0, 1)$, then we can use the following simple truncation argument. Notice that with probability $1 - o(n^{-100})$, a standard gaussian variable is bounded by $20\sqrt{\log n}$ (with room to spare). Thus, we can replace $N(0, 1)$ by its truncation at $20\sqrt{\log n}$, and set $K = 20\sqrt{\log n}$ and pay an extra term n^{-100} in the probability bound. One can apply this trick to any distribution with a light tail.

However, for some important applications, the entries of E can be reasonably large. A typical example here is the matrix completion problem, a fundamental problem in data science (see Section 4 for more details). In this case, E is not really noise in the traditional sense, but a random matrix created artificially from the setting of the problem. In this setting, the magnitude of K will make the error term in question become too big.

Example: Matrix completion. Let A be matrix with rank r and non-zero entries of order $O(1)$. Let B be a matrix obtained by keeping each entry of A with probability p , independently. We call these entries *observed*. For a non-observed entry, write 0. The task is to recreate A from B .

If we consider $\tilde{A} = \frac{1}{p}B$, then $\tilde{A}$ is a random matrix with expectation equalling A , thanks to the normalization. Thus, we can write $\tilde{A} = A + E$, where E is a random matrix with independent entries with mean 0. The entries ξ_{ij} of E have different distributions. For the ij entry, $\xi_{ij} = a_{ij}$ with probability $1 - p$, and $(1 - 1/p)a_{ij}$ with probability p . Thus K is roughly $\max_{ij} |a_{ij}|/p$ which is of order $1/p$. In the matrix completion problem, one often wants to make p as small as possible, typically $n^{-1+\epsilon}$ or even $\log n/n$. Thus, K can be close to n and the error in question becomes too big.

In this section, we develop new bounds to overcome this deficiency. We will need the notion of strong stability, which is a refinement of the notion of stability introduced earlier.

DEFINITION 3.2. We say that (σ_i, δ_i) is (c, τ, ν) strongly stable under E if in addition to being (c, τ, ν) stable, σ_i satisfies $\sigma_i > c\sqrt{Kn} \log^{\nu+0.01} n$.

ASSUMPTION 3.1. We assume for the rest of this section that $\mathbb{E}[\xi_{ij}^2] \leq K \leq n$.

This assumption is satisfied by random variables which take a large value K with a small probability of order $1/K$. This is exactly the situation with the matrix completion problem discussed above.

REMARK 3.4. For a random matrix of size n , whose entries have zero mean and variance K , the spectral norm is typically $\Omega(\sqrt{Kn})$. Thus, in this case, the last (new) condition in Definition 3.2 is only marginally stronger than the signal to noise assumption $\sigma_i \gg \|E\|$.

THEOREM 3.4. Let $c_0, \tau > 0$, with c_0 constant and τ potentially tending to 0 with n . Set $c = 2^{17}(c_0 + 1)r^3$. Assume that (σ_i, δ_i) is $(c, \tau, 2)$ strongly stable. Then with probability at least $1 - C(r)n^{-c_0} - \tau \log n$,

$$(3.10) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty \leq c \|U\|_\infty (\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + \kappa_i \epsilon_2(i) K \sqrt{\log n}) + \frac{c\sqrt{Kn} \kappa_i \|U\|_\infty \log n}{\sigma_i}.$$

REMARK 3.5. Results in random matrix theory show that $\|E\|$ often concentrates around $\sim \sqrt{Kn}$ [59]. In this case, the last term on the RHS of (3.10) is basically $\kappa_i \|U\|_\infty \epsilon_1(i)$. Thus, the bound essentially becomes

$$(3.11) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty = \tilde{O}\left(\kappa_i \|U\|_\infty (\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + K \epsilon_2(i))\right).$$

By the discussion in Remark 3.3, the $\epsilon_1(i)$ and $\epsilon_2(i)$ terms are necessary to bound $\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2$. The bound is thus

$$(3.12) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty = \tilde{O}\left(\kappa_i \|U\|_\infty \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2\right),$$

which is only off from (3.6) by a factor of κ_i . Thus, we basically achieve (3.6) even if K is large.

In a setting where $\|U\|_\infty = n^{-1/2+o(1)}$ and the condition number $\kappa_i = O(1)$, this new result essentially reduces K to $\sqrt{K}$. The bound on $\|U\|_\infty$ (incoherence bound) is typical and necessary in the analysis of the matrix completion problem; see for instance [18, 20, 44, 54].

The key step in our analysis, the so-called *delocalization* lemma below, is new and could be of independent interest.

LEMMA 3.1. (DELOCALIZATION LEMMA) Let $c_0, \tau > 0$, with c_0 constant. Set $c = 2^{11}(c_0 + 1)r^3$. If (σ_i, δ_i) is $(c, \tau, 2)$ strongly stable, then

$$(3.13) \quad \|\tilde{\mathbf{u}}_i\|_\infty \leq c_1 \kappa_i \|U\|_\infty,$$

with probability at least $1 - C(r)n^{-c_0} - \tau \log n$.

3.5 The rectangular case. The results are easy to generalize to the rectangular case, where $A \in \mathbb{R}^{m \times n}$, by a standard symmetrization trick. Let $N = m + n$. Define the $N \times N$ matrix $S(A) := \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}$. Notice that $S(A)$ is symmetric, and it is easy to show that if $A^T \mathbf{u} = \sigma \mathbf{v}$ and $A \mathbf{v} = \sigma \mathbf{u}$, then $S(A)(\mathbf{u}, -\mathbf{v}^T) = \sigma(\mathbf{u}, -\mathbf{v}^T)$. Thus, we can apply the main result to $S(A)$ to obtain ℓ_∞ perturbation bounds for $\mathbf{u}_i$ and $\mathbf{v}_i$, the i th left and right singular vectors of A . For this theorem, we will assume A has singular value decomposition $A = U\Sigma V$. We have the following rectangular analogue of Theorem 3.3

THEOREM 3.5. *Let $m_2(i) = \max\{\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2, \|\tilde{\mathbf{v}}_i - \mathbf{v}_i\|_2\}$, and let $W = [U, V]$, the concatenation of U and V . Let $c_0, \tau > 0$, with c_0 constant and τ potentially tending to 0 with n . Set $c = 2^{12}(c_0 + 1)r^3$. If (σ_i, δ_i) is $(c, \tau, 1)$ stable (with $\|W\|_\infty$ instead of $\|U\|_\infty$), we have with probability at least $1 - C(r)N^{-c_0} - 2\tau$,*

$$(3.14) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty \leq c \|U\|_\infty [\kappa_i m_2(i) + \epsilon_1(i) + \kappa_i \epsilon_2(i) K \sqrt{\log N}] + \frac{cK\sqrt{\log N}}{\sigma_i}$$

The same holds for $\tilde{\mathbf{v}}_i - \mathbf{v}_i$, with $\|U\|_\infty$ replaced with $\|V\|_\infty$.

Here is the analogue for Theorem 3.4

THEOREM 3.6. *Let $c_0, \tau > 0$, with c_0 constant and τ potentially tending to 0 with n . Set $c = 2^{18}(c_0 + 1)r^3$. If (σ_i, δ_i) is $(c, \tau, 2)$ strongly stable (replacing n in the definition of stable pair with N , and $\|U\|_\infty$ with $\|W\|_\infty$), then with probability at least $1 - C(r)N^{-c_0} - \tau \log N$,*

$$(3.15) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty \leq c \|U\|_\infty [\kappa_i m_2(i) + \epsilon_1(i) + \kappa_i \epsilon_2(i) K \sqrt{\log N}] + \frac{c\sqrt{KN}\kappa_i \|W\|_\infty \log N}{\sigma_i}.$$

The same holds for $\tilde{\mathbf{v}}_i - \mathbf{v}_i$, with $\|U\|_\infty$ replaced with $\|V\|_\infty$.

3.6 Sketch of the proofs and main new ideas. One can derive Theorem 3.3 from Theorem 3.2 by verifying that the assumptions of Theorem 3.2 hold with high probability in the setting of Theorem 3.3. This part requires a technical, but rather routine, computation.

In order to prove Theorem 3.2, we apply the leave-one-out strategy, which is a popular method to control the coordinates of an eigenvector. The starting observation here is the following. Let $H^{\{l\}}$ be the matrix obtained from H by zeroing out its l th row and column, then the l th row and column of $A + H^{\{l\}}$ and A are the same. On the other hand, the l th row of the matrix, thanks to the eigenvector equation $Av = \lambda v$, has a direct influence on the l th coordinate of any eigenvector v . From here, it is not hard to deduce a strong bound for the difference between the l th coordinates of an eigenvector of A and its counterpart of $A + H^{\{l\}}$.

By the triangle inequality, it remains to bound for the difference between the l th coordinates of the eigenvector of $A + H^{\{l\}}$ and its counterpart of $A + H$. It is often enough to just replace it by the ℓ_2 distance between the vectors. In many previous treatments, authors used the original Davis-Kahan to obtain a bootstrapping inequality [4, 25]. Our new idea here is to exploit the special structure of the difference matrix $H - H^{\{l\}}$, which has exactly one non-trivial row and column. Using a series of linear algebra manipulations, we obtain a more effective bound, laying the ground for a stronger bootstrap argument which results in the conclusion of Theorem 3.2.

To prove Theorem 3.4, we introduce the so-called *iterative leave-one-out* argument, which is a refinement of the original leave-one-out argument and could be of independent interest. The basic idea is as follows. Starting with the deterministic Theorem 3.2, we observe that quality of the bound provided by this theorem depends on the ℓ_∞ norm of the eigenvectors of the (leave-one-out) matrix $A + H^{\{l\}}$. To control this later quantity, we apply Theorem 3.2 again, but now on $A + H^{\{l\}}$. It will lead to a leave-two-out matrix, obtained by zeroing out two rows and columns of H . We keep continuing this process (leave-three-out and so on) and obtain a series of improvements, which converges to the desired bound. The key point here is that the more rows and columns we leave out, the weaker the requirements on the ℓ_∞ bound become, until a point that it is automatically satisfied.

Using this new argument, we first prove the Delocalization Lemma 3.1. It is then relatively simple to derive Theorem 3.4 from this lemma.

4 Applications

In this section, we apply our new results to a number of well known algorithmic problems, leading to fast and very simple algorithms.

4.1 Finding hidden partition. Finding hidden partition is a popular problem in statistics and theoretical computer science (also goes under the name of statistical block model). Here is the setting: a vertex set V of size n is partitioned into r subsets $V_1, \dots, V_r$, and between each pair V_i, V_j we draw edges independently with probability p_{ij} (we allow $i = j$). The task is to find a particular subset V_j or all the subsets $V_1, \dots, V_r$ given one instance of the random graph. See [7, 13, 28, 34, 38, 39, 40, 41, 43, 50, 61] and the references therein. We think of r as a constant and n tends to infinity.

The most popular approach to this problem is the spectral method (see [43] for a survey), which typically consists of two steps. In the first step, one considers the coordinates of an eigenvector of the adjacency matrix of the graph (or more generally the projection of the row vectors of the adjacency matrix onto a low dimensional eigenspace), and runs a standard clustering algorithm on these low dimensional n points. The output of this step is an approximation of the truth. In the second step, one applies adhoc combinatorial techniques to clean the output to recover the mis-classified vertices.

The input of the problem is the adjacency matrix of the (random) graph. Let us call this matrix $\tilde{A}$. Now let A be the matrix of expectation (thus the entries will be p_{ij}). Since there are r vertex sets in the partition, this matrix has r identical blocks and thus has rank at most r . The difference $E = \tilde{A} - A$ is a random matrix with independent upper diagonal entries. Since p_{ij} is the expectation of the ij entry of $\tilde{A}$, E has zero mean.

It has been speculated that in many cases, the cleaning step is not necessary. Our result makes a contribution towards solving this problem. The critical point here is that the existence of mis-classified vertices, in many settings, is just an artifact of the analysis in the first step, which typically relies on ℓ_2 norm estimates. It is clear that any ℓ_2 norm estimate, even sharp, could only imply that a majority of the vertices are well classified, which leads to the necessity of the second step. On the other hand, if we have a strong ℓ_∞ norm estimate, then we can classify all the vertices at once. Our new infinity norm estimates will enable us to do exactly this in a number of settings, resulting in simple and fast new algorithms. In Section [14], we apply this idea to many problems in this area, including the hidden clique problem, the planted coloring problem, the hidden bipartition problem, and the general hidden partition problem.

All of these problems have been studied heavily, with numerous treatments using different tools. On the other hand, our treatment is very simple and universal for all settings considered. Moreover, in certain ranges, the algorithm works under the weakest assumption known to date.

Let us close this section with an illustrative example.

The hidden clique problem. The (simplest form) of the hidden clique problem is the following: Hide a clique X of size k in the random graph $G(n, 1/2)$. Can we find X in polynomial time?

Notice that the largest clique in $G(n, 1/2)$, with overwhelming probability, has size approximately $2 \log n$ [8]. Thus, for any k bigger than $(2 + \epsilon) \log n$, with any constant $\epsilon > 0$, X would be abnormally large and therefore detectable, by brute-force at least. For instance, one can check all vertex sets of size k to see if any of them form a clique. However, finding X in polynomial time is a different matter, and the best current bound for k is $k \geq c\sqrt{n}$, for any constant $c > 0$. This was first achieved by Alon, Krivelevich, and Sudakov [7]; see also [40][30] for later developments concerning faster algorithms for certain values of c .

The Alon-Krivelevich-Sudakov algorithm runs as follows. It first finds X when c is sufficiently large, then uses a simple sampling trick to reduce the case of small c to this case.

To find the clique for a large c , they first compute the second eigenvector of the adjacency matrix of the graph and locate the first largest k coordinates in absolute value. Call this set Y . This is an approximation of the clique X , but not yet totally accurate. In the second, cleaning, step, they define X as the vertices in the graph with at least $3/4k$ neighbors in Y . The authors then proved that with high probability, X is indeed the hidden clique.

With our new results, we can find X immediately by a slightly modified version of the first step, omitting the cleaning step, as promised. Before starting the main step of the algorithm, we change all zeros in the adjacency matrix to -1 .

ALGORITHM 4.1. (FIRST SINGULAR VECTOR CLUSTERING-FSC) Compute the first singular vector. Let x be the

largest value of the coordinates and let X be the set of all coordinates with value at least $x/2$.

This is perhaps the simplest algorithm for this problem. Implementation is trivial as computing the first singular vector of a large matrix is a routine operation that appears in all standard numerical linear algebra packages.

THEOREM 4.1. *There is a constant c_0 such that for all $k \geq c_0\sqrt{n}$, FSC outputs the hidden clique correctly with probability at least .99.*

4.2 Matrix Completion. A major problem in data science is the matrix completion problem, which asks to recover a large matrix from a sparse, random, set of observed entries. Formally speaking, let A be a large $m \times n$ matrix where each entry is revealed with probability p , independently (thus roughly pmn entries are observed). The goal is to recover A from the set of observed entries.

One of the key motivations for this problem is to build rating/recommendation systems. Assume that a company wants to know customers' opinions about the entire catalog of their products. They can achieve this by constructing the rating matrix of their products, where the rows of represent customers and the columns are indexed by products, and each entry represents a rating. Clearly, entries of high ratings suggest a natural recommendation strategy.

The problem here is that only part of the matrix is known, as most customers have used and rated only few products. Thus, one needs to complete the matrix based on these few observed entries. A famous example here is the Netflix problem where the entries are the ratings of movies (from 1 to 5). In fact, matrix completion has become a public event thanks to the Netflix competition; see [47].

It is clear that the task is feasible only if there is some condition on the matrix, and the most popular condition is that A has low rank. There is a vast literature on the problem with this assumption; see [20, 20, 24, 54] and the references therein.

A natural try for matrix completion is to find the matrix of minimal rank agreeing with the observed entries. However, this problem is NP-hard. The idea here is to use the following relaxation

$$(4.16) \quad \text{minimize } \|X\|_* \text{ subject to. } X_{ij} = [P(A)]_{ij}, \text{ for all observed } (i, j)$$

where $\|X\|_*$ is the sum of the singular values of X . In words, the task is: among all matrices whose entries agree with the observed matrix entries, find the one with the smallest nuclear norm. A series of papers [18, 19, 20, 54], by Candès and many coauthors show that (under various assumptions) the solution to the convex program (4.16) recovers A exactly, with high probability.

Another idea is to use the spectral method. Consider a matrix $\tilde{A}$, where $\tilde{A}_{ij} = p^{-1}A_{ij}$ if the entry A_{ij} is observed, and 0 otherwise. Thus, $\tilde{A}_{ij}$ is a random variable with mean A_{ij} , as each entry is observed with probability p . So we can write $\tilde{A} = A + E$, where E is a random matrix with independent entries having zero mean. One can see $\tilde{A}$ as an unbiased estimator of A . A well known work in this direction is [44]. In this paper, Keshavan, Montanari, and Oh first use a low rank approximation of $\tilde{A}$ to obtain an approximation of A in Frobenius norm. Next, they solve an optimization problem to clean the output, and achieve exact recovery with high probability. See Table 1 for a summary of existing results.

In both approaches above, one needs to solve a non-trivial optimization problem. it seems that the nuclear norm optimization algorithm, while polynomial time in theory, does not run efficiently in practice; see [48].

As an application of our new results, we design a simple spectral algorithm, whose cleaning step is simply rounding the output of the spectral step. Assume for a moment that the entries, as in the Netflix problem, are non-zero integers. We are going to show that a properly chosen low rank approximation B of $\tilde{A}$ satisfies $\|A - B\|_\infty < 1/2$. Thus, one can recover A from B by simply rounding the entries to the nearest integer. This is thanks to the fact that we now can prove that a properly defined low rank approximation of $\tilde{A}$ approximates A in the *infinity norm*, compared to approximation in Frobenius norm or spectral norm in previous works.

We now describe the algorithm. First compute the leading singular values and singular vectors of $\tilde{A}$, $(\tilde{\sigma}_1, \tilde{u}_1, \tilde{v}_1), (\tilde{\sigma}_2, \tilde{u}_2, \tilde{v}_2), \dots, (\tilde{\sigma}_{\tilde{s}}, \tilde{u}_{\tilde{s}}, \tilde{v}_{\tilde{s}})$ where $\tilde{s} := \max_i \{i : \tilde{\sigma}_i \geq \frac{1}{8r} \|\tilde{W}\|_\infty^{-2}\}$. Let $s := \max_i \{i : \sigma_i \geq \frac{1}{16r} \|\tilde{W}\|_\infty^{-2}\}$. Observe that $\tilde{s}$ is random because it is computed from the observed $\tilde{A}$, while s is deterministic. B will be the low rank approximation given by

Result	Algorithm	Lowest possible density
Candés, Recht '09 [19]	convex optimization	$p = \Omega(N^{-0.8} \log N)$
Candés, Tao '10 [20]	convex optimization	$p = \Omega(N^{-1} \log^2 N)$
Recht '11 [54]	convex optimization	$p = \Omega(N^{-1} \log^2 n)$
Keshavan, Montanari, Oh '10 [44]	spectral + cleaning	$p = \Omega(N^{-1} \log N)$

Table 1: A survey of results for exact recovery in the matrix completion problem for a $m \times n$ matrix, and $N = m + n$. All results are stated under the assumptions that $r = O(1)$ and $\|U\|_\infty = O(n^{-1/2})$ in order to minimize the sampling density p .

$$B := \sum_{j=1}^{\tilde{s}} \tilde{\sigma}_j \tilde{u}_j \tilde{v}_j^T.$$

The formal code is as follows.

ALGORITHM 4.2. (APPROXIMATE-AND-ROUND)

1. Take SVD of $\tilde{A} = \sum_{i=1}^{\min\{m,n\}} \tilde{\sigma}_i \tilde{u}_i \tilde{v}_i^T$.
2. Approximate: let $B = \sum_{i \leq \tilde{s}} \tilde{\sigma}_i \tilde{u}_i \tilde{v}_i^T$
3. Round: round the entries of B to the nearest integer.

THEOREM 4.2. Let A be a $m \times n$ matrix of rank r whose entries are non-zero integers, where both $r, \|A\|_\infty = O(1)$. Let $N = m + n$ and $\bar{\delta} = \inf_{i \leq s} \delta_i$. Then, there exists $c = c(r, \|A\|_\infty)$ such that if

- (signal to noise) $\sigma_s > c(\sqrt{Np^{-1}}) \log^{2.01} N$
- (gap) $\bar{\delta} > cp^{-1} \log N$.
- (incoherence) $\|W\|_\infty \leq cN^{-1/2}$
- (density) $p > N^{-1} \log^{4.03} N$,

then Algorithm 4.2 recovers all of the entries of A exactly with probability at least $1 - N^{-1}$.

The optimal value for the density is $p = O(\log N/N)$, which has been essentially achieved in [44] (under various assumptions). In this paper, we focus on the simplicity of both the algorithm and the proof, so do not try to optimize p . A more sophisticated analysis will bring us close to the optimal bound, while keeping the algorithm essentially the same. This will be the topic of a future paper.

Finally, let us comment on our new assumptions. The assumption that the entries are integers is common for recommendation systems, as we have alluded to. Furthermore, in real life most data matrices become integral by multiplying by a relatively small constant. For instance, if all entries have at most 2 decimal places, then $100A$ is integral, and our algorithm works with an obvious re-scaling.

The assumption that the entries are non-zero is for convenience, and can be achieved by simply shifting the matrix. If we know that all entries are in the interval $[-L, L]$, for some integer $L > 0$, then $A + (L+1)J$ (where J is the all-one matrix) have non-zero entries in the interval $[1, 2L+1]$. Furthermore, the rank would change by at most 1. Thus, the shifted matrix basically has the same parameters as the original one.

4.3 Matrix completion with noise In a more realistic setting, many authors considered a model when the data matrix A is already corrupted by (random) noise, and we only observe a few entries from the corrupted matrix [4, 18, 26, 45].

This problem looks more technical than the original (noiseless) one. On the other hand, with respect to our approach, it is still exactly the same problem. Assume that each entry a_{ij} from A is corrupted by noise x_{ij} with mean zero. Thus, the corrupted matrix is $A' = A + X$. As argued before, the observed matrix is

$$(4.17) \quad \tilde{A} = A' + E',$$

where E' is a random matrix with independent entries ξ'_{ij} which are equal $(\frac{1}{p} - 1)(a_{ij} + x_{ij})$ with probability p , and $-(a_{ij} + x_{ij})$ with probability $1 - p$. Since x_{ij} are independent bounded random variables with mean zero, ξ'_{ij} are independent, zero mean and $O(1/p)$ bounded. This is still under the assumption of Theorem 4.2. Thus, we can easily deduce the following "noisy" version.

THEOREM 4.3. *Let A be a $m \times n$ matrix of rank r whose entries are non-zero integers, where both $r, \|A\|_\infty = O(1)$. For $1 \leq i \leq m, 1 \leq j \leq n$, let x_{ij} be B -bounded independent random variables with zero mean, and X be the matrix consisting of these entries. Suppose $B = O(1)$.*

Let $N = m + n$ and $\bar{\delta} = \inf_{i \leq s} \delta_i$. Then, there exists $c = c(r, \|A\|_\infty, B)$ such that if

- (signal to noise) $\sigma_s > c\sqrt{Np^{-1}} \log^{2.01} N$
- (gap) $\bar{\delta} > cp^{-1} \log N$.
- (incoherence) $\|W\|_\infty \leq cN^{-1/2}$
- (density) $p > N^{-1} \log^{4.03} N$,

then Algorithm 4.2 (given $A' + E'$ as input) recovers all of the entries of A exactly with probability at least $1 - N^{-1}$.

See Figure 1 for a numerical example.

The outline of the rest of the paper. We first collect preparatory tools from linear algebra and probability in Section 5. Then, we prove the deterministic Theorem 3.2 in Sections 6, 7, and 8. The proof of Theorem 3.3 from Theorem 3.2 is given in Section 9. Afterward, we present the proof of the delocalization lemma, Lemma 3.1, in Sections 10, 11, and 12. Next, we demonstrate in Section 13 that Theorem 3.4 follows fairly easily from the proof of Lemma 3.1. To conclude, we discuss in detail (with proofs) our applications in Sections 14 and 15.

5 Preparation

Throughout the paper, we will make repeated use of a few facts from linear algebra and probability. To make the exposition easier, we will collect these facts and some of their consequences here.


5.1 Linear Algebra.

FACT 5.1. (WEYL INEQUALITY) *Let A be a symmetric matrix with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ and singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$. Define $\tilde{A} = A + H$ for any symmetric matrix H . Assume that $\tilde{A}$ has eigenvalues and singular values $\tilde{\lambda}_i$ and $\tilde{\sigma}_i$, again ordered decreasingly. For all $1 \leq i \leq n$,*


$$(5.18) \quad \begin{aligned} \sigma_i - \|H\| &\leq \tilde{\sigma}_i \leq \sigma_i + \|H\|, \text{ and} \\ \lambda_i - \|H\| &\leq \tilde{\lambda}_i \leq \lambda_i + \|H\|. \end{aligned}$$

As an immediate consequence, we have the following two results.


1. *If $\sigma_i \geq 2\|H\|$, then $\tilde{\sigma}_i \geq \frac{\sigma_i}{2}$.*




(a) Original



(b) Corrupted



(c) Samples (of corrupted)



(d) Recovered

Figure 1: Approximate-and-Round run on $\tilde{A} = A' + E'$ as defined in (4.17), where A is a $N \times N$ block matrix of 2's and 4's and X is a ± 0.3 , mean zero, random matrix. Here, $N = 20,000$ and the sampling density is $p = 0.30$.

2. If $|\lambda_i| > \|H\|$, then $\tilde{\lambda}_i$ has the same sign as λ_i .

For any symmetric matrix H and index set α , let H^α be equal to the matrix obtained from H by zeroing out the rows and columns indexed by α (replacing all entries in these rows and columns by zeros). The following fact is well known and easy to prove.

FACT 5.2. For any index set α , $\|H^\alpha\| \leq \|H\|$.

5.2 Probability.

LEMMA 5.1. (Hoeffding's Inequality, Theorem 2.2.6 in [58]) Let $X_1, X_2, \dots, X_n$ be independent zero-mean random variables such that $a_i < X_i < b_i$ with probability 1. Then,

$$(5.19) \quad \mathbb{P}\left\{\left|\sum_{i=1}^n X_i\right| > t\right\} \leq 2 \exp\left(\frac{-2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right).$$

COROLLARY 5.1. Let $\mathbf{x}$ be a random vector whose entries are independent, zero-mean, K -bounded random variables. Then for any fixed unit vector $\mathbf{u}$ and any $C > 0$,

$$\mathbb{P}\{|\mathbf{x}^T \mathbf{u}| \geq CK\sqrt{\log n}\} \leq 2 \exp\left(-\frac{C^2}{2} \log n\right) = 2n^{-C^2/2}.$$

The same bound holds if $\mathbf{u}$ is a random unit vector from which $\mathbf{x}$ is independent.

LEMMA 5.2. (Bernstein's Inequality, Theorem 2.8.4 in [58]) Let $X_1, \dots, X_n$ be independent, K bounded, mean zero, random variables. Then

$$(5.20) \quad \mathbb{P}\left\{\left|\sum_{i=1}^n a_i X_i\right| \geq t\right\} \leq 2 \exp\left(-\frac{t^2/2}{\sum_{i=1}^n \mathbb{E}[X_i^2] + Kt/3}\right).$$

The following is a corollary of a result from [51]; see Appendix A for the proof.

THEOREM 5.1. Suppose that E is a symmetric random matrix with K -bounded, mean zero, independent entries above the diagonal. Suppose that A has rank r , and let $1 \leq k \leq r$ be an integer. Then, for any $t \geq 0$, the following hold.

$$(5.21) \quad \begin{aligned} &\mathbb{P}\{\tilde{\sigma}_k < \sigma_k - t\} \leq 4 \times 9^k \exp\left(-\frac{t^2}{128K^2}\right), \text{ and} \\ &\mathbb{P}\left\{\tilde{\sigma}_k > \sigma_k + t\sqrt{r} + 2\sqrt{k} \frac{\|E\|^2}{\tilde{\sigma}_k} + k \frac{\|E\|^3}{\tilde{\sigma}_k^2}\right\} \leq 4 \times 9^{2r} \exp\left(-r \frac{t^2}{128K^2}\right). \end{aligned}$$

In particular,

$$\mathbb{P}\left\{|\tilde{\sigma}_k - \sigma_k| > 2r\left(t + \frac{\|E\|^2}{\tilde{\sigma}_k} + \frac{\|E\|^3}{\tilde{\sigma}_k^2}\right)\right\} \leq 8 \times 9^{2r} \exp\left(-\frac{t^2}{128K^2}\right).$$

6 Proof of Theorem 3.2

In this proof, both A and H are deterministic. We will examine the effect of the full perturbation H on entry u_{il} by first considering the auxiliary perturbation $H^{\{l\}}$ (which is obtained from H by leaving out the l th row and column). This is an example of the so-called leave-one-out strategy, which has been used by many researchers in recent studies [4, 25, 33, 63, 64]. Next, we need to add the l th row and column back and consider the impact of these. This is the more technical part of the proof, which requires a careful analysis.

Let $H_{\{l\}} = H - H^{\{l\}}$. By definition, the entries outside the l th row and column of $H_{\{l\}}$ are all zero. We set $A^{\{l\}} := A + H^{\{l\}}$ and call the singular values and singular vectors of this matrix $\sigma_i^{\{l\}}$ and $\mathbf{u}_i^{\{l\}}$, respectively. The l th entry of $\mathbf{u}_i^{\{l\}}$ is $u_{il}^{\{l\}}$.

First, we show that the effect of $H^{\{l\}}$ on u_{il} is extremely small. This is the content of Lemma 6.1. Once this is established, we view $\tilde{A}$ as a perturbation of $A^{\{l\}}$, $\tilde{A} = A^{\{l\}} + H_{\{l\}}$. The structure of $H_{\{l\}}$ (now viewed as noise) will allow us to deduce a strong ℓ_2 bound for the leading singular vectors of $\tilde{A}$. The key here is that this bound will be so strong that even when we use it to upper bound the entry-wise perturbation, the result still leads to the claim of our theorem. This bound is the content of Lemma 6.2, which is the most technical part of the proof and requires some novel ideas, going far beyond applying the standard Davis-Kahan bound.

LEMMA 6.1. *Under the conditions of Theorem 3.2, for any $1 \leq l \leq n$,*

$$|u_{il}^{\{l\}} - u_{il}| \leq 2r \|U_{l,\cdot}\|_\infty \left[\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\|_2 + \epsilon_1(i) \right].$$

LEMMA 6.2. *Under the conditions of Theorem 3.2, for any $1 \leq l \leq n$,*

$$\|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\|_2 \leq 68r^{1/2} \left[[\epsilon_1(i) + a_l \epsilon_2(i)] |\tilde{u}_{il}| + \kappa_i a_l \|U_{l,\cdot}\|_\infty \epsilon_2(i) \right] + 32 \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i}.$$

Proof of Theorem 3.2 given the lemmas. By the triangle inequality and the fact that ℓ_2 norm dominates the ℓ_∞ norm, we have

$$|\tilde{u}_{il} - u_{il}| \leq |u_{il}^{\{l\}} - u_{il}| + \|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\|_2.$$

By Lemma 6.1, we have

$$\begin{aligned} |\tilde{u}_{il} - u_{il}| &\leq |u_{il}^{\{l\}} - u_{il}| + \|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\|_2 \\ &\leq 2r \|U_{l,\cdot}\|_\infty (\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\|_2 + \epsilon_1(i)) + \|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\|_2 \\ &\leq 2r \|U_{l,\cdot}\|_\infty (\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i)) + 4r \|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\|_2. \end{aligned}$$

Now using Lemma 6.2 to bound $\|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\|_2$, we obtain, with $C_0 = 4 * 272r^{3/2}$,

$$|\tilde{u}_{il} - u_{il}| \leq \frac{C_0}{4} \left[\kappa_i \|U_{l,\cdot}\|_\infty \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + [\epsilon_1(i) + a_l \kappa_i \epsilon_2(i)] \|U_{l,\cdot}\|_\infty + [\epsilon_1(i) + a_l \epsilon_2(i)] |\tilde{u}_{il}| \right] + 128r \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i}.$$

By the triangle inequality, we can bound $|\tilde{u}_{il}| \leq |u_{il}| + |\tilde{u}_{il} - u_{il}|$. This gives, letting $b = \epsilon_1(i) + a_l \epsilon_2(i)$,

$$|\tilde{u}_{il} - u_{il}| \leq \frac{C_0}{4} \left[\kappa_i \|U_{l,\cdot}\|_\infty \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + b \|U_{l,\cdot}\|_\infty + b(|u_{il}| + |\tilde{u}_{il} - u_{il}|) \right] + 128r \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i}.$$

Then, the coefficient in front of $|\tilde{u}_{il} - u_{il}|$ on the RHS is $\frac{C_0 b}{4}$. Recall that $\epsilon_1(i) = \frac{\|E\|}{\sigma_i}$ and $\epsilon_2(i) = \frac{1}{\delta_i}$. Therefore, because $\sigma_i \geq C_0 \|E\|$ and $\delta_i \geq C_0 a_l$, it must be the case that $b < \frac{2}{C_0}$. Then we can move all the terms with $|\tilde{u}_{il} - u_{il}|$ to the left with coefficient at most $1/2$. This gives

$$(6.22) \quad \frac{1}{2} |u_{ij} - \tilde{u}_{ij}| \leq \frac{C_0}{4} \left[\kappa_i \|U_{l,\cdot}\|_\infty \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + b \|U_{l,\cdot}\|_\infty + b |u_{il}| \right] + 128r \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i}.$$

Bounding the term $|u_{il}|$ on the RHS by $\|U_{l,\cdot}\|_\infty$, we obtain

$$\frac{1}{2} |\tilde{u}_{il} - u_{il}| \leq \frac{C_0 \|U_{l,\cdot}\|_\infty}{4} \left[\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + 2b \right] + 128r \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i}.$$

Multiplying both sides by 2, we obtain the desired inequality (with room to spare).

□

In the next two sections, we establish Lemmas 6.1 and 6.2 respectively.

7 Proof of Lemma 6.1

Recall that for a matrix M , $M_{l,\cdot}$ denotes the l th row of M , viewed as a vector. The key point of the leave-one-out analysis is that by definition, $A_{l,\cdot}^{\{l\}} = A_{l,\cdot}$. Furthermore, as $\mathbf{u}_i^{\{l\}}$ is a singular vector of $A^{\{l\}}$, $A^{\{l\}}\mathbf{u}_i^{\{l\}}$ is either $\sigma_i^{\{l\}}\mathbf{u}_i^{\{l\}}$ or $-\sigma_i^{\{l\}}\mathbf{u}_i^{\{l\}}$ and we will use the shorthand $A^{\{l\}}\mathbf{u}_i^{\{l\}} = \pm\sigma_i^{\{l\}}\mathbf{u}_i^{\{l\}}$. In all estimates where this shorthand appears, the sign does not matter. Since $A^{\{l\}}\mathbf{u}_i^{\{l\}} = \pm\sigma_i^{\{l\}}\mathbf{u}_i^{\{l\}}$,

$$u_{il}^{\{l\}} = \frac{\langle A_{l,\cdot}^{\{l\}}, \mathbf{u}_i^{\{l\}} \rangle}{\pm\sigma_i^{\{l\}}} = \frac{\langle A_{l,\cdot}, \mathbf{u}_i^{\{l\}} \rangle}{\pm\sigma_i^{\{l\}}}.$$

By the spectral decomposition of A , the RHS can be written as

$$\frac{\langle A_{l,\cdot}, \mathbf{u}_i^{\{l\}} \rangle}{\pm\sigma_i^{\{l\}}} = \frac{1}{\pm\sigma_i^{\{l\}}} \left[\pm\sigma_1 u_{1l} \mathbf{u}_1^T \pm\sigma_2 u_{2l} \mathbf{u}_2^T + \cdots \pm\sigma_r u_{rl} \mathbf{u}_r^T \right] \mathbf{u}_i^{\{l\}},$$

which implies, via the triangle inequality, that

$$|u_{il}^{\{l\}} - u_{il}| \leq \left| \frac{\sigma_i}{\sigma_i^{\{l\}}} u_{il} \mathbf{u}_i^T \mathbf{u}_i^{\{l\}} - u_{il} \right| + \left| \frac{1}{\sigma_i^{\{l\}}} \sum_{j \neq i}^r \pm\sigma_j u_{jl} \mathbf{u}_j^T \mathbf{u}_i^{\{l\}} \right|.$$

In the first term on the RHS, we have eliminated the signs in front of σ_i and $\sigma_i^{\{l\}}$. This is because their signs correspond to the signs of the corresponding eigenvalues of A and $A^{\{l\}}$ respectively. It must be the case that these eigenvalues have the same sign by Fact 5.1. Studying the second term on the RHS, write $\mathbf{u}_i^{\{l\}} = \mathbf{u}_i + (\mathbf{u}_i^{\{l\}} - \mathbf{u}_i)$. By the orthogonality of $\mathbf{u}_i$ with $\mathbf{u}_j$ for $j \neq i$, it follows that

$$(7.23) \quad |u_{il}^{\{l\}} - u_{il}| \leq \left| \frac{\sigma_i}{\sigma_i^{\{l\}}} u_{il} \mathbf{u}_i^T \mathbf{u}_i^{\{l\}} - u_{il} \right| + \frac{2}{\sigma_i} \sum_{j \neq i}^r \sigma_j |u_{jl}| \left\| \mathbf{u}_i^{\{l\}} - \mathbf{u}_i \right\|_2.$$

To bound first term on the right hand side of (7.23), write

$$(7.24) \quad \frac{\sigma_i}{\sigma_i^{\{l\}}} u_{il} \mathbf{u}_i^T \mathbf{u}_i^{\{l\}} - u_{il} = u_{il} \left(\frac{\sigma_i - \sigma_i^{\{l\}}}{\sigma_i^{\{l\}}} \right) \mathbf{u}_i^T \mathbf{u}_i^{\{l\}} + u_{il} \mathbf{u}_i^T (\mathbf{u}_i^{\{l\}} - \mathbf{u}_i).$$

Since $\sigma_i^{\{l\}} > \sigma_i/2$ by Facts 5.1 and 5.2, the triangle inequality gives

$$(7.25) \quad \begin{aligned} \left| \frac{\sigma_i}{\sigma_i^{\{l\}}} u_{il} \mathbf{u}_i^T \mathbf{u}_i^{\{l\}} - u_{il} \right| &\leq |u_{il}| \frac{|\sigma_i - \sigma_i^{\{l\}}|}{\sigma_i^{\{l\}}} \left\| \mathbf{u}_i \right\|_2 \left\| \mathbf{u}_i^{\{l\}} \right\|_2 + |u_{il}| \left\| \mathbf{u}_i \right\|_2 \left\| \mathbf{u}_i^{\{l\}} - \mathbf{u}_i \right\|_2 \\ &\leq 2|u_{il}| \frac{\|H\|}{\sigma_i} + |u_{il}| \left\| \mathbf{u}_i^{\{l\}} - \mathbf{u}_i \right\|_2 \\ &\leq 2\epsilon_1(i) \|U_{l,\cdot}\|_\infty + \left\| \mathbf{u}_i^{\{l\}} - \mathbf{u}_i \right\|_2 \|U_{l,\cdot}\|_\infty. \end{aligned}$$

The second line uses Fact 5.1 to bound $|\sigma_i - \sigma_i^{\{l\}}| \leq \|H^{\{l\}}\|$, and Fact 5.2 to get $\|\mathbf{u}_i^{\{l\}}\| \leq \|H\|$. To bound the second term on the RHS of (7.23), bound $|u_{il}| \leq \|U_{l,\cdot}\|_\infty$ to obtain

$$\frac{2}{\sigma_i} \sum_{j \neq i}^r \sigma_j |u_{jl}| \left\| \mathbf{u}_i^{\{l\}} - \mathbf{u}_i \right\|_2 \leq 2 \|U_{l,\cdot}\|_\infty \frac{\left\| \mathbf{u}_i^{\{l\}} - \mathbf{u}_i \right\|_2 \sigma_1 (r-1)}{\sigma_i}.$$

We can bound this last term by $2\kappa_i \|U_{l,\cdot}\|_\infty \left\| \mathbf{u}_i^{\{l\}} - \mathbf{u}_i \right\|_2 (r-1)$ to conclude that

$$\begin{aligned}
 |u_{il}^{\{l\}} - u_{il}| &\leq 2\epsilon_1(i) \|U_{l,\cdot}\|_\infty + \left\| \mathbf{u}_i^{\{l\}} - \mathbf{u}_i \right\|_2 \|U_{l,\cdot}\|_\infty + 2\kappa_i \|U_{l,\cdot}\|_\infty \left\| \mathbf{u}_i^{\{l\}} - \mathbf{u}_i \right\|_2 (r-1) \\
 (7.26) \quad &\leq 2r\kappa_i \|U_{l,\cdot}\|_\infty \left\| \mathbf{u}_i^{\{l\}} - \mathbf{u}_i \right\|_2 + 2\epsilon_1(i) \|U_{l,\cdot}\|_\infty \\
 &\leq 2r \left[\kappa_i \|U_{l,\cdot}\|_\infty \left\| \tilde{\mathbf{u}}_i - \mathbf{u}_i \right\|_2 + \kappa_i \|U_{l,\cdot}\|_\infty \left\| \tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}} \right\|_2 + \epsilon_1(i) \|U_{l,\cdot}\|_\infty \right].
 \end{aligned}$$

where the last step uses the triangle inequality. This concludes the proof of Lemma 6.1

8 Proof of Lemma 6.2

In this section, we will view $\tilde{A}$ as a perturbation of $A^{\{l\}}$ with the perturbing matrix $H_{\{l\}}$. The main idea is that $H_{\{l\}}$ is only supported on one row and one column. By leveraging this and $\|U_{l,\cdot}\|_\infty$, we can obtain a strong bound for $\left\| \tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}} \right\|_2$. We begin with the following decomposition which will prove useful throughout:

$$(8.27) \quad H_{\{l\}} = \mathbf{x} \mathbf{e}_l^T + \mathbf{e}_l \mathbf{x}^T.$$

Recall that $\mathbf{x}$ is the l th row of H , but with the l th entry set to $H_{ll}/2$, as $H_{\{l\}}$ is $H - H^{\{l\}}$. Define $p = \min\{j : \sigma_{j+1} < \sigma_i/4\}$. Let $P^{\{l\}}$ be the orthogonal projection to the orthogonal complement of the columns of $U_p^{\{l\}}$, and let $V_p^{\{l\}}$ be the $n \times (p-1)$ matrix whose columns are $\mathbf{u}_1^{\{l\}}, \dots, \mathbf{u}_{i-1}^{\{l\}}, \mathbf{u}_{i+1}^{\{l\}}, \dots, \mathbf{u}_p^{\{l\}}$. Expanding $\tilde{\mathbf{u}}_i$ in the coordinates of the orthonormal basis $\{\mathbf{u}_k^{\{l\}}\}_{1 \leq k \leq n}$,

$$\tilde{\mathbf{u}}_i = \sum_{k=1}^p \alpha_k \mathbf{u}_k^{\{l\}} + P^{\{l\}} \tilde{\mathbf{u}}_i.$$

It follows that

$$\begin{aligned}
 \left\| \tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}} \right\|_2^2 &= \langle \tilde{\mathbf{u}}_i, \tilde{\mathbf{u}}_i \rangle + \langle \mathbf{u}_i^{\{l\}}, \mathbf{u}_i^{\{l\}} \rangle - 2\langle \tilde{\mathbf{u}}_i, \mathbf{u}_i^{\{l\}} \rangle \\
 &= 2(1 - \alpha_i^2) \\
 (8.28) \quad &= 2 \sum_{k \neq i}^n \alpha_k^2 \\
 &= 2 \sum_{k \neq i}^p \alpha_k^2 + 2 \left\| P^{\{l\}} \tilde{\mathbf{u}}_i \right\|_2^2 \\
 &= 2 \left\| V_p^{\{l\}T} \tilde{\mathbf{u}}_i \right\|_2^2 + 2 \left\| P^{\{l\}} \tilde{\mathbf{u}}_i \right\|_2^2.
 \end{aligned}$$

Therefore,

$$(8.29) \quad \left\| \tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}} \right\|_2 \leq \sqrt{2} \left[\left\| V_p^{\{l\}T} \tilde{\mathbf{u}}_i \right\|_2 + \left\| P^{\{l\}} \tilde{\mathbf{u}}_i \right\|_2 \right].$$

Proving Lemma 6.2 reduces to bounding the two terms on the RHS. It is possible that the first term in the fourth line of (8.28) is a sum over an empty set (say $i = 1, p = 1$). In this case, $\left\| \tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}} \right\|_2 \leq \sqrt{2} \left\| P^{\{l\}} \tilde{\mathbf{u}}_i \right\|_2$.

LEMMA 8.1. ($\left\| P^{\{l\}} \tilde{\mathbf{u}}_i \right\|_2$ BOUND)

$$(8.30) \quad \left\| P^{\{l\}} \tilde{\mathbf{u}}_i \right\|_2 \leq 24\epsilon_1(i)(|\tilde{u}_{il}| + \left\| \tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}} \right\|_2) + 8 \frac{|\langle \mathbf{x}, \mathbf{u}_i^{\{l\}} \rangle|}{\sigma_i}.$$

As we previously observed, there is no contribution from $\left\|V_p^{\{l\}T} \tilde{\mathbf{u}}_i\right\|_2$ when the aforementioned sum is empty. So we assume without loss of generality that it is not. We first establish that it is sufficient to bound the quantity $\left\|U_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i\right\|_2$ by using the perturbation technique of [51]. This is the content of the following proposition, which is where we use the gap stability condition.

PROPOSITION 8.1.

$$(8.31) \quad \left\|V_p^{\{l\}T} \tilde{\mathbf{u}}_i\right\|_2 \leq 2\epsilon_2(i) \left\|U_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i\right\|_2.$$

Having established this, when we go to bound $\left\|U_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i\right\|_2$, the structure of $H_{\{l\}}$ will bring the l th row of $U_p^{\{l\}}$ into play. It is important that this row not be too large in norm.

PROPOSITION 8.2. (THE l TH ROW OF $U_p^{\{l\}}$ IS SMALL) *Let $\mathbf{r}^{\{l\}}$ denote the l th row of $U_p^{\{l\}}$ viewed as a column vector. Then*

$$(8.32) \quad \left\|\mathbf{r}^{\{l\}}\right\|_2 \leq 8r^{1/2} \kappa_i \|U_{l,\cdot}\|_\infty.$$

The two propositions can be shown to give us a bound for $\left\|V_p^{\{l\}T} \tilde{\mathbf{u}}_i\right\|_2$.

LEMMA 8.2.

$$(8.33) \quad \left\|V_p^{\{l\}T} \tilde{\mathbf{u}}_i\right\|_2 \leq 16r^{1/2} \epsilon_2(i) \left[a_l(|\tilde{u}_{il}| + \kappa_i \|U_{l,\cdot}\|_\infty) + \kappa_i \|H\| \|U_{l,\cdot}\|_\infty \left\|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\right\|_2 \right].$$

We are now ready to prove Lemma 6.2

Proof of Lemma 6.2 given Lemmas 8.1 and 8.2. Lemmas 8.1 and 8.2 can be used to bound the RHS of (8.29). In particular, temporarily setting $\beta := \left\|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\right\|_2$ for brevity,

$$(8.34) \quad \begin{aligned} \beta &\leq \sqrt{2} \left[\left\|V_p^{\{l\}T} \tilde{\mathbf{u}}_i\right\|_2 + \left\|P^{\{l\}} \tilde{\mathbf{u}}_i\right\|_2 \right] \\ &\leq 34r^{1/2} \left[\epsilon_1(i) |\tilde{u}_{il}| + a_l \epsilon_2(i) (|\tilde{u}_{il}| + \kappa_i \|U_{l,\cdot}\|_\infty) + [\epsilon_1(i) + \kappa_i \|H\| \|U_{l,\cdot}\|_\infty \epsilon_2(i)] \beta \right] + 16 \frac{|\langle \mathbf{x}, \mathbf{u}_i^{\{l\}} \rangle|}{\sigma_i}. \end{aligned}$$

The main observation is that β appears on both the LHS and RHS of the inequality. The coefficient of β in the RHS is $34r^{1/2} [\epsilon_1(i) + \kappa_i \|H\| \|U_{l,\cdot}\|_\infty \epsilon_2(i)]$. By the definition of $\epsilon_1(i)$ and $\epsilon_2(i)$ (see the discussion preceding Theorem 3.2), this equals $34r^{1/2} (\frac{\|H\|}{\sigma_i} + \kappa_i \|H\| \|U_{l,\cdot}\|_\infty)$. By the definition of C_0 , and the assumption that $\sigma_i > C_0 \|E\|$ and $\delta_i > C_0 \kappa_i \|H\| \|U\|_\infty$, we have the following estimate.

$$(8.35) \quad \beta \leq 34r^{1/2} \left[\epsilon_1(i) |\tilde{u}_{il}| + a_l \epsilon_2(i) (|\tilde{u}_{il}| + \kappa_i \|U_{l,\cdot}\|_\infty) \right] + \frac{1}{2} \beta + 16 \frac{|\langle \mathbf{x}, \mathbf{u}_i^{\{l\}} \rangle|}{\sigma_i}.$$

Therefore, moving the terms involving β to the left and multiplying both sides by 2 gives

$$\left\|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\right\|_2 \leq 68r^{1/2} \left[[\epsilon_1(i) + a_l \epsilon_2(i)] |\tilde{u}_{il}| + \kappa_i a_l \|U_{l,\cdot}\|_\infty \epsilon_2(i) \right] + 32 \frac{|\langle \mathbf{x}, \mathbf{u}_i^{\{l\}} \rangle|}{\sigma_i}.$$

□

Proof of Lemma 8.1. Since $\tilde{A} - A^{\{l\}} = H_{\{l\}}$, it follows that

$$(8.36) \quad (P^{\{l\}}\tilde{\mathbf{u}}_i)^T \tilde{A}\tilde{\mathbf{u}}_i - (P^{\{l\}}\tilde{\mathbf{u}}_i)^T A^{\{l\}}\tilde{\mathbf{u}}_i = (P^{\{l\}}\tilde{\mathbf{u}}_i)^T H_{\{l\}}\tilde{\mathbf{u}}_i.$$

By the definition of $P^{\{l\}}$, we have

$$(8.37) \quad |(P^{\{l\}}\tilde{\mathbf{u}}_i)^T A^{\{l\}}\tilde{\mathbf{u}}_i| = |\langle P^{\{l\}}\tilde{\mathbf{u}}_i, A^{\{l\}}\tilde{\mathbf{u}}_i \rangle| = |\langle \tilde{\mathbf{u}}_i, P^{\{l\}} A^{\{l\}}\tilde{\mathbf{u}}_i \rangle| \leq \sigma_{p+1}^{\{l\}} \|P^{\{l\}}\tilde{\mathbf{u}}_i\|_2^2.$$

By Fact 5.1, $\sigma_{p+1}^{\{l\}} \leq \sigma_{p+1} + \|H^{\{l\}}\|$. Furthermore, by Fact 5.2, $\|H^{\{l\}}\| \leq \|H\|$, so we have $\sigma_{p+1}^{\{l\}} \leq \sigma_{p+1} + \|H\|$. Because $\tilde{\mathbf{u}}_i$ is a singular vector of $\tilde{A}$, we have

$$(P^{\{l\}}\tilde{\mathbf{u}}_i)^T \tilde{A}\tilde{\mathbf{u}}_i = \pm \tilde{\sigma}_i \|P^{\{l\}}\tilde{\mathbf{u}}_i\|_2^2.$$

It thus follows that

$$(8.38) \quad \tilde{\sigma}_i \|P^{\{l\}}\tilde{\mathbf{u}}_i\|_2^2 - (\sigma_{p+1} + \|H\|) \|P^{\{l\}}\tilde{\mathbf{u}}_i\|_2^2 \leq |(P^{\{l\}}\tilde{\mathbf{u}}_i)^T H_{\{l\}}\tilde{\mathbf{u}}_i|.$$

Applying Cauchy-Schwarz on the RHS, we obtain

$$(8.39) \quad \tilde{\sigma}_i \|P^{\{l\}}\tilde{\mathbf{u}}_i\|_2^2 - (\sigma_{p+1} + \|H\|) \|P^{\{l\}}\tilde{\mathbf{u}}_i\|_2^2 \leq \|P^{\{l\}}\tilde{\mathbf{u}}_i\|_2 \|H_{\{l\}}\tilde{\mathbf{u}}_i\|_2.$$

By Fact 5.1, $\tilde{\sigma}_i(A) > \frac{\sigma_i}{2}$. By definition of p , $\frac{\sigma_i}{2} - \sigma_{p+1} > \frac{1}{4}\sigma_i$. So dividing by $\|P^{\{l\}}\tilde{\mathbf{u}}_i\|_2$ gives

$$(8.40) \quad \|P^{\{l\}}\tilde{\mathbf{u}}_i\|_2 \leq \frac{\|H_{\{l\}}\tilde{\mathbf{u}}_i\|_2}{0.25\sigma_i - \|H\|} \leq \frac{\|H_{\{l\}}\mathbf{u}_i^{\{l\}}\|_2 + \|H_{\{l\}}\| \|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\|_2}{0.25\sigma_i - \|H\|} \leq 8 \frac{\|H_{\{l\}}\mathbf{u}_i^{\{l\}}\|_2 + 2\|H\| \|\tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}}\|_2}{\sigma_i}.$$

We used the triangle inequality in the numerator. It is apparent that $\mathbf{x}$ has ℓ_2 norm at most that of the l th row of H . This gave $\|H_{\{l\}}\| = \|\mathbf{x}\mathbf{e}_l^T + \mathbf{e}_l\mathbf{x}^T\| \leq 2\|\mathbf{x}\| \leq 2\|H\|$. We also lower bounded $0.25\sigma_i - \|H\| \geq \sigma_i/8$ because $\sigma_i > C_0\|H\|$. To estimate the term $\|H_{\{l\}}\mathbf{u}_i^{\{l\}}\|_2$, write using 8.27,

$$(8.41) \quad H_{\{l\}}\mathbf{u}_i^{\{l\}} = \langle \mathbf{x}, \mathbf{u}_i^{\{l\}} \rangle \mathbf{e}_l + u_{il}^{\{l\}} \mathbf{x}.$$

Here $\mathbf{e}_l$ is the l th standard basis vector. We obtain

$$(8.42) \quad \begin{aligned} \|H_{\{l\}}\mathbf{u}_i^{\{l\}}\|_2 &\leq |\langle \mathbf{x}, \mathbf{u}_i^{\{l\}} \rangle| + \|\mathbf{x}\|_2 |u_{il}^{\{l\}}| \\ &\leq |\langle \mathbf{x}, \mathbf{u}_i^{\{l\}} \rangle| + \|H\| (|\tilde{u}_{il}| + |\tilde{u}_{il} - u_{il}^{\{l\}}|). \end{aligned}$$

So we conclude that

$$(8.43) \quad \|P^{\{l\}}\tilde{\mathbf{u}}_i\|_2 \leq \frac{8|\langle \mathbf{x}, \mathbf{u}_i^{\{l\}} \rangle| + 24\|H\| (|\tilde{u}_{il}| + \|\mathbf{u}_i^{\{l\}} - \tilde{\mathbf{u}}_i\|_2)}{\sigma_i},$$

proving the lemma. $\square$

Proof of Proposition 8.1. By the relation $\tilde{A} - A^{\{l\}} = H_{\{l\}}$, we have

$$(8.44) \quad V_p^{\{l\}T} \tilde{A} \tilde{\mathbf{u}}_i - V_p^{\{l\}T} A^{\{l\}} \tilde{\mathbf{u}}_i = V_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i.$$

First, we observe that because $\tilde{\mathbf{u}}_i$ is a singular vector of $\tilde{A}$, the first term on the LHS of (8.44) is $\pm \tilde{\sigma}_i V_p^{\{l\}T} \tilde{\mathbf{u}}_i$. Since the columns of $V_p^{\{l\}}$ are singular vectors of $A^{\{l\}}$, the second term on the left hand side of (8.44) is $D^{\{l\}} V_p^{\{l\}T} \tilde{\mathbf{u}}_i$. $D^{\{l\}}$ is the $(p-1) \times (p-1)$ diagonal matrix with entries $\pm \sigma_1^{\{l\}}, \dots, \pm \sigma_{i-1}^{\{l\}}, \pm \sigma_{i+1}^{\{l\}}, \dots, \pm \sigma_p^{\{l\}}$. Then, we have

$$(8.45) \quad \left\| U_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i \right\|_2 \geq \left\| V_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i \right\|_2 = \left\| (\pm \tilde{\sigma}_i I - D^{\{l\}}) V_p^{\{l\}T} \tilde{\mathbf{u}}_i \right\|_2 \geq \frac{\delta_i}{2} \left\| V_p^{\{l\}T} \tilde{\mathbf{u}}_i \right\|_2.$$

The first inequality uses that $V_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i$ is a sub-vector of $U_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i$. The equality uses (8.44), and the last inequality uses that the smallest singular value of the diagonal matrix $\pm \tilde{\sigma}_i I - D^{\{l\}}$ is at least $\delta_i/2$ by the assumption of Theorem 3.2.

□

Proof of Proposition 8.2. Since the columns of $U_p^{\{l\}}$ are singular vectors, if $\Sigma_p^{\{l\}}$ is a diagonal matrix with entries $\pm \sigma_1^{\{l\}}, \dots, \pm \sigma_p^{\{l\}}$, we have by definition of $\mathbf{r}^{\{l\}}$,

$$(8.46) \quad \left\| \mathbf{r}^{\{l\}} \right\|_2 = \left\| \mathbf{e}_l^T U_p^{\{l\}} \right\|_2 = \left\| \mathbf{e}_l^T A^{\{l\}} U_p^{\{l\}} \Sigma_p^{\{l\}-1} \right\|_2.$$

Because $A^{\{l\}}$ and A have the same l th row,

$$(8.47) \quad \left\| \mathbf{r}^{\{l\}} \right\|_2 = \left\| \mathbf{e}_l^T A U_p^{\{l\}} \Sigma_p^{\{l\}-1} \right\|_2.$$

By the spectral decomposition of A , if Σ is the diagonal matrix whose entries are the eigenvalues of A , then the RHS is $\left\| \mathbf{e}_l^T U \Sigma U^T U_p^{\{l\}} \Sigma_p^{\{l\}-1} \right\|_2 = \left\| U_{l,\cdot} \Sigma U^T U_p^{\{l\}} \Sigma_p^{\{l\}-1} \right\|_2$. By the triangle inequality,

$$(8.48) \quad \left\| \mathbf{r}^{\{l\}} \right\|_2 \leq \|U_{l,\cdot}\|_2 \|\Sigma\| \left\| U^T U_p^{\{l\}} \right\| \left\| \Sigma_p^{\{l\}-1} \right\|_2.$$

Because $\sigma_1 > C_0 \|E\|$ and the definition of p , $\Sigma_p^{\{l\}-1}$ has norm at most $\sigma_p^{\{l\}} \leq \sigma_p/2$ by Facts 5.1 and 5.2. It is clear that $\|\Sigma\| = \sigma_1$ and $\left\| U^T U_p^{\{l\}} \right\| \leq 1$. Therefore, since $\|U_{l,\cdot}\|_2 \leq \sqrt{r} \|U_{l,\cdot}\|_\infty$,

$$(8.49) \quad \left\| \mathbf{r}^{\{l\}} \right\|_2 \leq 2r^{1/2} \left(\frac{\sigma_1}{\sigma_p} \right) \|U_{l,\cdot}\|_\infty \leq 8\kappa_i r^{1/2} \|U_{l,\cdot}\|_\infty,$$

where the last inequality uses the definition of p .

□

Proof of Lemma 8.2. By Proposition 8.1, it suffices to upper bound $\left\| U_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i \right\|_2$. Using the decomposition for $H_{\{l\}}$ (8.27), we can write

$$U_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i = (H_{\{l\}} U_p^{\{l\}})^T \tilde{\mathbf{u}}_i = (X + Y) \tilde{\mathbf{u}}_i.$$

Thus, we must bound $\|X \tilde{\mathbf{u}}_i\|_2 + \|Y \tilde{\mathbf{u}}_i\|_2$, where $X = U_p^{\{l\}T} \mathbf{x} \mathbf{e}_l^T$ and $Y = U_p^{\{l\}T} \mathbf{e}_l \mathbf{x}^T$. Therefore,

$$(8.50) \quad \|X \tilde{\mathbf{u}}_i\|_2 \leq |\tilde{u}_{il}| \left\| U_p^{\{l\}T} \mathbf{x} \right\|_2 = |\tilde{u}_{il}| a_l.$$

On the other hand, for $\|Y \tilde{\mathbf{u}}_i\|_2$, the triangle inequality gives

$$(8.51) \quad \|Y \tilde{\mathbf{u}}_i\|_2 \leq \left\| Y \mathbf{u}_i^{\{l\}} \right\|_2 + \|Y\| \left\| \tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}} \right\|_2.$$

By definition of Y , we can write $Y \mathbf{u}_i^{\{l\}} = \langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle U_p^{\{l\}T} \mathbf{e}_l = \langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle \mathbf{r}^{\{l\}}$. As before, $\mathbf{e}_l$ is the l th standard basis vector. Invoking Proposition 8.2 it follows immediately that

$$(8.52) \quad \left\| Y \mathbf{u}_i^{\{l\}} \right\|_2 \leq |\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle| \left\| \mathbf{r}^{\{l\}} \right\|_2 \leq 8r^{1/2} \kappa_i a_l \|U_{l,\cdot}\|_\infty.$$

Now, we need to deal with the term $\|Y\| \left\| \tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}} \right\|_2$ in (8.51). We will bound $\|Y\|$. It is easy to see that Y can be written as $Y = \mathbf{r}^{\{l\}} \mathbf{x}^T$. In particular, Y is rank 1 so we can calculate the spectral norm of Y directly as $\|Y\| = \|\mathbf{x}\|_2 \left\| \mathbf{r}^{\{l\}} \right\|_2$. We have previously observed that $\|\mathbf{x}\|_2 \leq \|H\|$ and we can bound $\left\| \mathbf{r}^{\{l\}} \right\|_2$ as we did before using Proposition 8.2. Therefore,

$$(8.53) \quad \|Y\| \leq 8r^{1/2} \kappa_i \|H\| \|U_{l,\cdot}\|_\infty.$$

Combining estimates (8.50), (8.51), (8.52), and (8.53) gives

$$\begin{aligned} \left\| U_p^{\{l\}T} H_{\{l\}} \tilde{\mathbf{u}}_i \right\|_2 &\leq \|X \tilde{\mathbf{u}}_i\|_2 + \|Y \tilde{\mathbf{u}}_i\|_2 \\ &\leq |\tilde{u}_{il}| a_l + 8r^{1/2} \kappa_i a_l \|U_{l,\cdot}\|_\infty + 8r^{1/2} \kappa_i \|H\| \|U_{l,\cdot}\|_\infty \left\| \tilde{\mathbf{u}}_i - \mathbf{u}_i^{\{l\}} \right\|_2. \end{aligned}$$

□

9 Proof of Theorem 3.3 via Theorem 3.2

In this section, we deduce Theorem 3.3 from the deterministic Theorem 3.2. The task is basically checking that the conditions of Theorem 3.2 hold with high probability. The hardest part is the stability condition, and for this, we will need to appeal to singular value perturbation bounds from [51]. We will show that on the complement of a bad event $\mathcal{B}$, the conditions for Theorem 3.2 hold for $H = E$. Next, the bad event holds with small probability.

Define the event $\mathcal{B} := \mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_E$ where

$$(9.54) \quad \begin{aligned} \mathcal{B}_1 &:= \cup_{1 \leq l \leq n} \{ \min\{|\tilde{\sigma}_i - \sigma_{i+1}^{\{l\}}|, |\tilde{\sigma}_i - \sigma_{i-1}^{\{l\}}|\} < \delta_i/2 \}, \\ \mathcal{B}_2 &:= \cup_{1 \leq l \leq n} \left\{ \left\| U^{\{l\}T} \mathbf{x}(l) \right\|_2 \geq \sqrt{2r(c_0 + 1) \log n} \right\}, \text{ and} \\ \mathcal{B}_E &:= \{ \|E\| > T \}. \end{aligned}$$

LEMMA 9.1.

$$(9.55) \quad \mathbb{P}(\mathcal{B}) \leq C(r)n^{-c_0} + 2\tau.$$

Proof of Theorem 3.3 given Theorem 3.2 and Lemma 9.1. By the construction of $\mathcal{B}$, we can check that on $\overline{\mathcal{B}}$ (complement of $\mathcal{B}$), the conditions for Theorem 3.2 with $H = E$ and $H^{\{l\}} = E^{\{l\}}$ are satisfied for all l .

Verification of the conditions for Theorem 3.2. Let $1 \leq l \leq n$. We first have to check that on $\overline{\mathcal{B}}$,

$$\min\{|\tilde{\sigma}_i - \sigma_{i+1}^{\{l\}}|, |\tilde{\sigma}_i - \sigma_{i-1}^{\{l\}}|\} \geq \delta_i/2.$$

The definition of $\mathcal{B}_1$ makes this condition trivial. Then, we have to check that

$$\sigma_i > C_0 \|E\|.$$

On $\overline{\mathcal{B}}$, the event $\mathcal{B}_E$ guarantees that $\|E\| \leq T$. Therefore, by $(c, \tau, 1)$ stability,

$$\sigma_i > cT \geq c\|E\| > C_0 \|E\|.$$

See condition (a) of Definition 3.1. Finally, we verify that

$$(9.56) \quad \delta_i > C_0 \max \left\{ \left\| U^{\{l\}T} \mathbf{x}(l) \right\|_2, \kappa_i \|E\| \|U\|_\infty \right\}.$$

On $\bar{\mathcal{B}}$, the event $\mathcal{B}_2$ guarantees that

$$\left\| U^{\{l\}T} \mathbf{x}(l) \right\|_2 \leq K \sqrt{2r(c_0 + 1) \log n}.$$

Because of the $(c, \tau, 1)$ stability assumption,

$$\delta_i > c \max\{K \sqrt{\log n}, T \kappa_i \|U\|_\infty\}.$$

Since $\|E\| \leq T$, this ensures that (9.56) holds. See conditions (b) and (c) of Definition 3.1

Conclusion of the proof. By Lemma 9.1, $\mathcal{B}$ has probability at most $C(r)n^{-c_0} + 2\tau$ (recall that $C(r) = 1000 \times 9^{2r}$). Furthermore, we have checked that, if $\bar{\mathcal{B}}$ occurs, then for all $1 \leq l \leq n$, the conditions for Theorem 3.2 hold. For each l , Theorem 3.2 gives that

$$(9.57) \quad |\tilde{u}_{il} - u_{il}| \leq C_0 \|U_{l,\cdot}\|_\infty \left[\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + a_l \kappa_i \epsilon_2(i) \right] + 256r \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i},$$

where we recall $a_l = \|U^{\{l\}T} \mathbf{x}\|_2$. Taking the maximum over l on both sides gives us that on $\bar{\mathcal{B}}$,

$$(9.58) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty \leq C_0 \|U\|_\infty \left[\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + (\max_l a_l) \kappa_i \epsilon_2(i) \right] + 256r \frac{\max_l |\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i}.$$

Notice that $\max_l |\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle| \leq \max_l a_l$. Then we can use the bound for $\max_l a_l \leq \sqrt{2r(c_0 + 1) \log n}$ on $\bar{\mathcal{B}}_2$ to conclude that

$$(9.59) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty \leq c \|U\|_\infty \left[\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + (K \sqrt{\log n}) \kappa_i \epsilon_2(i) \right] + c \frac{K \sqrt{\log n}}{\sigma_i},$$

which is precisely (3.9). $\square$

Thus, what remains is to prove Lemma 9.1.

9.1 Proof of Lemma 9.1. We will bound $\mathbb{P}(\mathcal{B}_E)$, $\mathbb{P}(\mathcal{B}_2)$, and $\mathbb{P}(\mathcal{B}_1)$ separately, and use the union bound.

Probability of $\mathcal{B}_E$. Recall that

$$\mathcal{B}_E = \{\|E\| > T\}.$$

By the definition of T , it must be the case that

$$(9.60) \quad \mathbb{P}(\mathcal{B}_E) \leq \tau.$$

Probability of $\mathcal{B}_2$. Recall that

$$\mathcal{B}_2 = \cup_{1 \leq l \leq n} \left\{ \left\| U^{\{l\}T} \mathbf{x} \right\|_2 \geq \sqrt{2r(c_0 + 1) \log n} \right\}.$$

Let $1 \leq l \leq n$. Observe that $\|U^{\{l\}T} \mathbf{x}\|_2$ is the length of the projection of a random vector onto a subspace from which it is independent. Consider the vector $U^{\{l\}T} \mathbf{x}$, which has r entries. We use Corollary 5.1 to bound each entry $[U^{\{l\}T} \mathbf{x}]_j$ of this vector. We can bound for any $1 \leq j \leq r$,

$$(9.61) \quad \mathbb{P}\{|[U^{\{l\}T} \mathbf{x}]_j| \geq K \sqrt{2(c_0 + 1) \log n}\} \leq 2n^{-c_0-1}.$$

By taking the union bound over the r entries,

$$(9.62) \quad \mathbb{P}\left\{\left\|U^{\{l\}T}\mathbf{x}\right\|_2 \geq K\sqrt{2r(c_0+1)\log n}\right\} \leq 2rn^{-c_0-1}.$$

Thus $\mathcal{B}_2$ holds with probability at most $2rn^{-c_0}$ by taking a union bound over $1 \leq l \leq n$.

$$(9.63) \quad \mathbb{P}(\mathcal{B}_2) \leq 2rn^{-c_0}.$$

Probability of $\mathcal{B}_1$. What remains is to bound the probability of $\mathcal{B}_1$. This is the hardest step, so we treat it separately.

LEMMA 9.2.

$$(9.64) \quad \mathbb{P}(\mathcal{B}_1) \leq 64 \times 9^{2r}n^{-c_0} + \tau.$$

By using the union bound, (9.60), (9.63), and (9.64) imply Lemma 9.1. We dedicate the remainder of the section to proving Lemma 9.2, which completes the proof of Lemma 9.1.

9.2 Proof of Lemma 9.2. Define

$$(9.65) \quad \begin{aligned} \mathcal{G}_{i-1} &:= \bigcap_{1 \leq l \leq n} \left\{ \max_{k=i-1, i} \max\{|\sigma_k^{\{l\}} - \sigma_k|, |\tilde{\sigma}_k - \sigma_k|\} \leq 24r \left[K\sqrt{r(c_0+1)\log n} + \frac{\|E\|^2}{\tilde{\sigma}_k} + \frac{\|E\|^3}{\tilde{\sigma}_k^2} \right] \right\}, \\ \mathcal{G}_{i+1} &:= \bigcap_{1 \leq l \leq n} \left\{ \max_{k=i, i+1} \max\{|\sigma_k^{\{l\}} - \sigma_k|, |\tilde{\sigma}_k - \sigma_k|\} \leq 24r \left[K\sqrt{r(c_0+1)\log n} + \frac{\|E\|^2}{\tilde{\sigma}_k} + \frac{\|E\|^3}{\tilde{\sigma}_k^2} \right] \right\}. \end{aligned}$$

These good events essentially guarantee that the relevant perturbed singular values are close to the original ones.

PROPOSITION 9.1.

$$(9.66) \quad \mathbb{P}(\mathcal{B}_1) \leq \mathbb{P}(\overline{\mathcal{G}_{i-1}}) + \mathbb{P}(\overline{\mathcal{G}_{i+1}}) + \mathbb{P}(\mathcal{B}_E).$$

Proof of Lemma 9.2 given the proposition. We apply the results of [51] on the perturbation of singular values to determine the probability of $\overline{\mathcal{G}_{i+1}}$ and $\overline{\mathcal{G}_{i-1}}$. This is because for all l , both E and $E^{\{l\}}$ satisfy the conditions of Theorem 5.1. The bound obtained from Theorem 5.1 for both E and $E^{\{l\}}$ will be the same because $\|E^{\{l\}}\| \leq \|E\|$ by Fact 5.2.

We will bound $\mathbb{P}(\overline{\mathcal{G}_{i+1}})$, and the exact same bound will hold for $\mathbb{P}(\overline{\mathcal{G}_{i-1}})$. Let $1 \leq l \leq n$. We apply Theorem 5.1 to $k \in \{i, i+1\}$ with $t = K\sqrt{128(c_0+1)\log n}$. For each such k ,

$$(9.67) \quad \mathbb{P}\left\{\max\{|\sigma_k^{\{l\}} - \sigma_k|, |\tilde{\sigma}_k - \sigma_k|\} \geq K\sqrt{128r(c_0+1)} + 2\sqrt{r}\frac{\|E\|^2}{\tilde{\sigma}_k} + r\frac{\|E\|^3}{\tilde{\sigma}_k^2}\right\} \leq 16 \times 9^{2r} \exp[-(c_0+1)\log n].$$

This implies that

$$(9.68) \quad \mathbb{P}\left\{\max_{k=i, i+1} \max\{|\sigma_k^{\{l\}} - \sigma_k|, |\tilde{\sigma}_k - \sigma_k|\} \leq 24r \left[K\sqrt{r(c_0+1)\log n} + \frac{\|E\|^2}{\tilde{\sigma}_k} + \frac{\|E\|^3}{\tilde{\sigma}_k^2} \right] \right\} \leq 32 \times 9^{2r}n^{-c_0-1}.$$

Since $\mathcal{G}_{i+1}$ is an intersection over l , we have to take a union bound over $1 \leq l \leq n$ to bound the complement. Then, $\overline{\mathcal{G}_{i+1}}$ holds with probability at most $32 \times 9^{2r}n^{-c_0}$. The same bound holds for $\overline{\mathcal{G}_{i-1}}$.

Recall that in (9.60), we already bounded $\mathbb{P}(\mathcal{B}_E)$. Putting the bounds for $\overline{\mathcal{G}_{i-1}}$, $\overline{\mathcal{G}_{i+1}}$, and $\mathbb{P}(\mathcal{B}_E)$ together, Proposition 9.1 implies

$$(9.69) \quad \mathbb{P}(\mathcal{B}_1) \leq 64 \times 9^{2r}n^{-c_0} + \tau.$$

□

Proof of Proposition 9.1. Recall that

$$\mathcal{B}_1 = \cup_{1 \leq l \leq n} \min\{|\tilde{\sigma}_i - \sigma_{i+1}^{\{l\}}|, |\tilde{\sigma}_i - \sigma_{i-1}^{\{l\}}|\} \leq \delta_i/2\}, \text{ and}$$

$$\mathcal{B}_E = \{\|E\| \geq T\}.$$

Let

$$\mathcal{B}_{1,i-1} = \cup_{1 \leq l \leq n} \{|\tilde{\sigma}_i - \sigma_{i-1}^{\{l\}}| \leq \delta_i/2\}, \text{ and}$$

$$\mathcal{B}_{1,i+1} = \cup_{1 \leq l \leq n} \{|\tilde{\sigma}_i - \sigma_{i+1}^{\{l\}}| \leq \delta_i/2\}.$$

Observe that $\mathcal{B}_1 \subset \mathcal{B}_{1,i-1} \cup \mathcal{B}_{1,i+1}$. We will show that

$$(9.70) \quad \mathcal{B}_{i,i-1} \subset \overline{\mathcal{G}_{i-1}} \cup \mathcal{B}_E, \text{ and}$$

$$(9.71) \quad \mathcal{B}_{i,i+1} \subset \overline{\mathcal{G}_{i+1}} \cup \mathcal{B}_E.$$

Then, the union bound will imply the proposition. Recall that $\Delta_i = \sigma_i - \sigma_{i+1}$, and $\delta_i = \min(\Delta_{i-1}, \Delta_i)$. Since the proof of (9.70) is virtually identical, we will only give the proof of (9.71). We break the analysis up into cases depending on if $\Delta_i > 4T$ or not.

Case 1. $\Delta_i > 4T$.

Suppose $\overline{\mathcal{B}_E}$ holds. Let $1 \leq l \leq n$. By Fact 5.2, $\|E^{\{l\}}\| \leq \|E\|$, and on $\overline{\mathcal{B}_E}$, $\|E\| \leq T$. Therefore, by applying Weyl's inequality to $\sigma_{i+1}^{\{l\}}$ and $\tilde{\sigma}_i$, we find that

$$|\tilde{\sigma}_i - \sigma_{i+1}^{\{l\}}| \geq \sigma_i - \sigma_{i+1} - 2T > \Delta_i/2 \geq \delta_i/2.$$

Therefore, $\overline{\mathcal{B}_E} \subset \overline{\mathcal{B}_{1,i+1}}$, so this implies that $\mathcal{B}_{1,i+1} \subset \mathcal{B}_E \subset \mathcal{B}_E \cup \overline{\mathcal{G}_{i+1}}$ in this case.

Case 2. $\Delta_i \leq 4T$.

We will show that $\overline{\mathcal{G}_{i+1}} \cup \mathcal{B}_E = \mathcal{G}_{i+1} \cap \overline{\mathcal{B}_E} \subset \overline{\mathcal{B}_{i,i+1}}$, implying (9.71). Suppose $\mathcal{G}_{i+1} \cap \overline{\mathcal{B}_E}$ holds. Let $1 \leq l \leq n$. Since $\sigma_i > 100T$ by $(c, \tau, 1)$ stability (and choice of c) and $\Delta_i \leq 4T$,

$$\sigma_{i+1} = \sigma_i - \Delta_i \geq \frac{9}{10}\sigma_i.$$

We can use Weyl's inequality to obtain

$$\tilde{\sigma}_{i+1} \geq \sigma_{i+1} - \|E\| \geq \frac{9}{10}\sigma_i - T \geq \frac{8}{10}\sigma_i.$$

With this lower bound for $\tilde{\sigma}_{i+1}$, we can upper bound the right hand side of the inequality defining $\mathcal{G}_{i+1}$ for both $k = i$ and $k = i + 1$. On $\mathcal{G}_{i+1} \cap \overline{\mathcal{B}_E}$, we have that

$$(9.72) \quad \forall k \in \{i+1, i\} : \max\{|\sigma_k^{\{l\}} - \sigma_k|, |\tilde{\sigma}_k - \sigma_k|\} \leq 32r \left[K\sqrt{r(c_0+1)\log n} + \frac{T^2}{\sigma_i} + \frac{T^3}{\sigma_i^2} \right].$$

Since $\sigma_i > 100T$, the third term on the right hand side of the above inequality is at most $1/100$ of the second term. This implies that on $\mathcal{G}_{i+1} \cap \overline{\mathcal{B}_E}$,

$$(9.73) \quad \forall k \in \{i+1, i\} : \max\{|\sigma_k^{\{l\}} - \sigma_k|, |\tilde{\sigma}_k - \sigma_k|\} \leq 33r \left[K\sqrt{r(c_0+1)\log n} + \frac{T^2}{\sigma_i} \right].$$

By the definition of c , it must be the case that Δ_i is much larger than the RHS, because $\Delta_i > c(K\log^{1/2} n + \sigma_i^{-1}T^2)$ by assumption (b) of $(c, \tau, 1)$ stability.

$$(9.74) \quad \forall k \in \{i+1, i\} : \max\{|\sigma_k^{\{l\}} - \sigma_k|, |\tilde{\sigma}_k - \sigma_k|\} \leq \frac{\Delta_i}{4}.$$

Thus, on $\mathcal{G}_{i+1} \cap \overline{\mathcal{B}_E}$, both $\tilde{\sigma}_i$ and $\sigma_{i+1}^{\{l\}}$ are at most $\frac{\Delta_i}{4}$ away from their original values (which are σ_i and σ_{i+1}), so the gap between them is at least $\Delta_i/2 \geq \delta_i/2$. Since this is true for all l , this implies that $\overline{\mathcal{B}_{i,i+1}}$ holds. Thus, in either case, $\mathcal{B}_{1,i+1} \subset \overline{\mathcal{G}_{i+1}} \cup \mathcal{B}_E$.

As we have stated, a virtually identical argument using a case analysis for Δ_{i-1} gives that $\mathcal{B}_{1,i-1} \subset \overline{\mathcal{G}_{i-1}} \cup \mathcal{B}_E$. Therefore,

$$\mathcal{B}_1 \subset \overline{\mathcal{G}_{i-1}} \cup \overline{\mathcal{G}_{i+1}} \cup \mathcal{B}_E,$$

which implies that

$$(9.75) \quad \mathbb{P}(\mathcal{B}_1) \leq \mathbb{P}(\overline{\mathcal{G}_{i-1}}) + \mathbb{P}(\overline{\mathcal{G}_{i+1}}) + \mathbb{P}(\mathcal{B}_E).$$

□

10 Proof of the Delocalization Lemma 3.1

We now prove Lemma 3.1, using the iterative leave-one-out argument, discussed briefly in Section 3.6.

Recall the bound from Theorem 3.2 with $A, H = E$, on coordinate l :

$$(10.76) \quad |\tilde{u}_{il} - u_{il}| \leq C_0 \|U_{l,\cdot}\|_\infty \left[\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + a_l \kappa_i \epsilon_2(i) \right] + 256r \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i},$$

where we recall that C_0 is the value from Theorem 3.2, $C_0 = 272 \times 4r^{3/2}$. The term in the brackets will be shown to be less than $\frac{5}{2}\kappa_i$. Further, $\|U_{l,\cdot}\|_\infty \leq \|U\|_\infty$. We can thus write

$$(10.77) \quad \begin{aligned} |\tilde{u}_{il} - u_{il}| &\leq \frac{5}{2} C_0 \kappa_i \|U\|_\infty + 256r \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i}, \text{ so} \\ |\tilde{u}_{il}| &\leq 3C_0 \kappa_i \|U\|_\infty + 256r \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x} \rangle|}{\sigma_i}. \end{aligned}$$

The last term on the RHS is the inner product of the singular vector of a leave-one-out matrix with a random vector from which it is independent. Applying the Hoeffding inequality here is somewhat wasteful, as we can apply the stronger Bernstein inequality, given that we have control on the infinity norm of $\mathbf{u}_i^{\{l\}}$. Thus, the problem reduces to bounding the infinity norm of eigenvectors of a minor. On the surface, this makes the problem harder, as there are n minors. But we observe that the bound for the minor is slightly weaker than what we need for the whole matrix. This gain is critical and we are able to exploit it in a full iterative argument. The details now follow.

Notation. Let α be an index set. Set E^α to be the random matrix equal to E , but with the rows and columns indexed by α set to zero. This is a generalization of the leave-one-out construction from Theorem 3.2. If $|\alpha| = j$, we have a leave- j -out matrix. Let $A^\alpha = A + E^\alpha$. U^α will be the matrix of r leading singular vectors of A^α . Similar notations, such as $\mathbf{u}_i^\alpha$ and σ_i^α , are self-explanatory.

Roughly speaking, the iterative leave-one-out argument uses a weaker bound on $\max_{|\alpha|=j} \|\mathbf{u}_i^\alpha\|_\infty$ to obtain a stronger bound for $\max_{|\alpha|=j-1} \|\mathbf{u}_i^\alpha\|_\infty$. We will define a deterministic, increasing sequence $\{f_j\}_{0 \leq j \leq j^*}$, where j^* is a number smaller than $\log n$. The sequence will be defined so that that $f_j^* = 1$, and f_1 and f_0 will be of size at most $4C_0 \kappa_i \|U\|_\infty$. We will show that under the $(c, \tau, 2)$ strong stability assumption, the leave- j -out singular vectors satisfy the following for $0 \leq j \leq j^*$.

$$(10.78) \quad \text{For all } |\alpha| = j, \|\mathbf{u}_i^\alpha\|_\infty \leq f_j.$$

Bounding the probability of failure of this statement is tricky. It involves conditioning on what happens at step j (leaving- j -out) to control what happens at step $j-1$ (leaving- $(j-1)$ -out). This iterative bound is proven in Lemma 10.1. The low probability of failure of (10.78) for $j=0$ will conclude the proof of Lemma 3.1.

We now formally define our parameters. For $0 \leq j \leq j^* := \lceil \frac{50 \log n}{\log(\log n)} + 3 \rceil$, define f_j in the following fashion. Start with $f_{j^*} = 1$. For $0 \leq j \leq j^* - 1$,

$$(10.79) \quad f_j := 3C_0\kappa_i \|U\|_\infty + \frac{f_{j+1}}{\log^{0.01} n}.$$

By the choice of j^* and the fact that $C_0 \geq 1$, it is easy to check that f_0 and f_1 are both less than $4C_0\kappa_i \|U\|_\infty$.

LEMMA 10.1. (ITERATIVE LEMMA) For $0 \leq j \leq j^*$, define

$$\gamma_j := \mathbb{P}\left\{\max_{\alpha, |\alpha|=j} \|\mathbf{u}_i^\alpha\|_\infty > f_j\right\}.$$

Then $\gamma_{j^*} = 0$. Further, for $1 \leq j \leq j^*$,

$$\gamma_{j-1} \leq \gamma_j + 2\tau + \epsilon_j,$$

with $\epsilon_j = 66 \times 9^{2r} n^j \exp(-c_2 \log^2 n)$, and $c_2 = c_0 + 1$.

The first conclusion that $\gamma_{j^*} = 0$ is trivial, as a coordinate of a unit vector is at most 1, and we defined $f_{j^*} = 1$. The important content of this lemma is thus the iterative bound for the γ_j .

Proof of Lemma 3.1 given Lemma 10.1 We have the relation

$$\sum_{j=1}^{j^*-1} \epsilon_j \leq n\epsilon_{j^*-1} = \epsilon_{j^*}.$$

Lemma 10.1 implies that

$$(10.80) \quad \gamma_0 \leq 2j^*\tau + \epsilon_{j^*} + \sum_{j=1}^{j^*-1} \epsilon_j \leq 2j^*\tau + 2\epsilon_{j^*}.$$

Restating this with the definition of γ_0 and ϵ_{j^*} ,

$$(10.81) \quad \begin{aligned} \mathbb{P}\{\|\tilde{\mathbf{u}}_i\|_\infty > f_0\} &\leq 2j^*\tau + 132 \times 9^{2r} \times n^{j^*} \exp(-c_2 \log^2 n) \\ &= 2j^*\tau + 132 \times 9^{2r} \times \exp((j^* - c_2 \log n) \log n) \\ &\leq \tau \log n + 132 \times 9^{2r} \exp(-\omega(\log n)) \\ &= \tau \log n + 132 \times 9^{2r} n^{-\omega(1)} \\ &\leq \tau \log n + 132 \times 9^{2r} n^{-c_0} \end{aligned}$$

for any constant $c_0 > 0$, thanks to the fact that $j^* = O(\log n / \log \log n) = o(\log n)$. We have previously observed that $f_0 \leq 4C_0\kappa_i \|U\|_\infty$; therefore (10.81) implies Lemma 3.1 $\square$

We now prove Lemma 10.1

Preliminaries. Recall that for an index l , we defined $\mathbf{x}(l)$ as the l th row of E with the l th entry divided by 2. We now let $\mathbf{x}(\alpha, l)$ be the l th row of E^α with its l th entry divided by 2. By definition, any entry of $\mathbf{x}(\alpha, l)$ is either zero, an entry of E , or an entry of E divided by 2. In particular, the entries of $\mathbf{x}(\alpha, l)$ are mean zero, K -bounded, independent random variables. We consider this vector because in the deterministic Theorem 3.2 the l th row of H plays an important role when we bound the perturbation of the l th coordinate of $\mathbf{u}_i$. We will apply Theorem 3.2 with $H = E^\alpha$.

The proof of Lemma 10.1 requires bounding the probability of various failure events, which we now define.

The event $\mathcal{B}_{\alpha, l}$.

Let α be an index set. Let $1 \leq l \leq n$, and set $\beta = \alpha \cup \{l\}$. Let

$$\begin{aligned}
 \mathcal{B}_{\alpha,l} &:= \mathcal{B}_{\alpha,l,1} \cup \mathcal{B}_{\alpha,l,2} \cup \mathcal{B}_{\alpha,E}, \text{ where} \\
 \mathcal{B}_{\alpha,l,1} &:= \{\min\{|\sigma_i^\alpha - \sigma_{i+1}^\beta|, |\sigma_i^\alpha - \sigma_{i-1}^\beta|\} < \delta_i/2\}, \\
 \mathcal{B}_{\alpha,l,2} &:= \left\{ \|U^{\beta T} \mathbf{x}(\alpha, l)\|_2 \geq K \sqrt{2c_2 r \log^2 n} \right\}, \\
 \mathcal{B}_{\alpha,E} &:= \{\|E^\alpha\| > T\}.
 \end{aligned}
 \tag{10.82}$$

For $0 \leq j \leq j^*$, let

$$\mathcal{B}_j := \bigcup_{\substack{|\alpha|=j \\ l \notin \alpha}} (\mathcal{B}_{\alpha,l,1} \cup \mathcal{B}_{\alpha,l,2}) \cup \bigcup_{|\alpha|=j} \mathcal{B}_{\alpha,E}.$$

LEMMA 10.2. (PROBABILITY OF $\mathcal{B}_j$) *Let $0 \leq j \leq j^*$. Under the conditions of Lemma [3.1](#),*

$$\mathbb{P}(\mathcal{B}_j) \leq 65 \times 9^{2r} n^{j+1} \exp(-c_2 \log^2 n) + 2\tau.$$

The event $\mathcal{F}_{\alpha,l}$.

For an index set α with $|\alpha| = j$, and a coordinate $1 \leq l \leq n$, define

$$\mathcal{F}_{\alpha,l} := \{|u_{il}^\alpha| > f_j\}.$$

We wish to show that the entries of $\mathbf{u}_i^\alpha$ are small. When $\mathcal{F}_{\alpha,l}$ holds, it means that a coordinate of $\mathbf{u}_i^\alpha$ is too big, and represents a failure at level j .

We will need the following lemma about the $\mathcal{F}_{\alpha,l}$ for those α, l where $l \in \alpha$.

LEMMA 10.3.

$$\bigcup_{\substack{|\alpha|=j \\ l \in \alpha}} \mathcal{F}_{\alpha,l} \subset \bigcup_{|\alpha|=j} \mathcal{B}_{\alpha,E}.
 \tag{10.83}$$

Having discussed how to control $\mathcal{F}_{\alpha,l}$ when $l \in \alpha$, we move to the case where $l \notin \alpha$. The remaining events are for controlling the probability of $\mathcal{F}_{\alpha,l}$ for such α, l .

The event $\mathcal{K}_{\alpha,l}$.

For an index set α with $|\alpha| = j$, and a coordinate $1 \leq l \leq n$, set $\beta = \alpha \cup \{l\}$. Define

$$\mathcal{K}_{\alpha,l} := \left\{ |u_{il}^\alpha| > 3C_0 \kappa_i \|U\|_\infty + 256r \frac{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle|}{\sigma_i} \right\}.
 \tag{10.84}$$

If the success event $\overline{\mathcal{K}_{\alpha,l}}$ occurs, we can show that $|u_{il}^\alpha|$ is small provided the inner product $\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle$ is small. The following lemma shows that $\mathcal{K}_{\alpha,l}$ is unlikely.

LEMMA 10.4. ($\mathcal{K}_{\alpha,l} \subset \mathcal{B}_{\alpha,l}$ WHEN $l \notin \alpha$.) *Let $0 \leq j \leq j^*$. Let α be an index set such that $|\alpha| = j$, and let $l \notin \alpha$. Then,*

$$\mathcal{K}_{\alpha,l} \subset \mathcal{B}_{\alpha,l}.
 \tag{10.85}$$

If $l \notin \alpha$, $|\beta| = j + 1$. Showing that the inner product $\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle$ is small (thus showing that $|u_{il}^\alpha|$ is small) requires information about the infinity norm of $\mathbf{u}_i^\beta$. This information will be provided by the following events.

The events $\mathcal{L}_{\alpha,l}$ and $\mathcal{I}_{\alpha,l}$.

For an index set α with $|\alpha| = j$, and a $1 \leq l \leq n$, let $\beta = \alpha \cup \{l\}$. Define

$$\begin{aligned}
 \mathcal{L}_{\alpha,l} &:= \{\|\mathbf{u}_i^\beta\|_\infty > f_{j+1}\}, \\
 \mathcal{I}_{\alpha,l} &:= \{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle| \geq c_2 \sqrt{2Kn} f_{j+1} \log^2 n\}, \text{ and} \\
 \mathcal{I}_j &:= \bigcup_{\substack{|\alpha|=j \\ l \notin \alpha}} (\mathcal{I}_{\alpha,l} \cap \overline{\mathcal{L}_{\alpha,l}}).
 \end{aligned}$$

LEMMA 10.5. (PROBABILITY OF $\mathcal{I}_j$) Let $0 \leq j \leq j^*$. Under the conditions of Lemma 3.1,

$$(10.86) \quad \mathbb{P}(\mathcal{I}_j) \leq 2n^{j+1} \exp(-c_2 \log^2 n).$$

Proof of Lemma 10.1 given the lemmas. We have previously observed that the statement is trivially true for $j = j^*$ because $\gamma_{j^*}^* = 0$. Having handled this, we move to the proof of the iterative bound. Let $1 \leq j \leq j^*$. Consider a set α with $j-1$ elements, which defines matrices $E^\alpha, A^\alpha = A + E^\alpha$. Let l be a coordinate such that $l \notin \alpha$, and let $\beta = \alpha \cup \{l\}$. We aim to apply Theorem 3.2 for the pair A, E^α on coordinate l with E^α playing the role of H . The theorem obtains a bound on $|u_{il}^\alpha - u_{il}|$. This bound implies that

$$(10.87) \quad |u_{il}^\alpha| \leq 3C_0\kappa_i \|U\|_\infty + 256r \frac{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle|}{\sigma_i}.$$

Looking at the second term on the RHS, if

$$(10.88) \quad |\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle| \leq c_2 \sqrt{2Kn} f_j \log^2 n,$$

then the RHS of (10.87) is at most f_{j-1} . This is because by the $(c, \tau, 2)$ strong stability assumption, $\sigma_i > c\sqrt{Kn} \log^{2.01} n$. Since $c > 256rc_2\sqrt{2}$, (10.87) and (10.88) imply that

$$(10.89) \quad \begin{aligned} |u_{il}^\alpha| &\leq 3C_0\kappa_i \|U\|_\infty + 256r \frac{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle|}{\sigma_i} \\ &\leq 3C_0\kappa_i \|U\|_\infty + f_j \frac{256rc_2\sqrt{2Kn} \log^2 n}{\sigma_i} \\ &\leq 3C_0\kappa_i \|U\|_\infty + \frac{f_j}{\log^{0.01} n} \\ &= f_{j-1}, \end{aligned}$$

and we thus have

$$(10.90) \quad |u_{il}^\alpha| \leq f_{j-1}.$$

The occurrence of (10.90) is the event we are interested in. Recall that $\mathcal{F}_{\alpha,l}$ is the event where (10.90) fails. What we have shown is that for those α, l such that $l \notin \alpha$, the failure probability of (10.90) and thus the bound for $\mathcal{F}_{\alpha,l}$ consists of two components. The first component is that the inequality (10.87) does not hold, which we recall is the failure event $\mathcal{K}_{\alpha,l}$. By Lemma 10.4, $\mathcal{K}_{\alpha,l} \subset \mathcal{B}_{\alpha,l}$. The second component is that the inner product bound (10.88) fails. We called this failure event $\mathcal{I}_{\alpha,l}$. To analyze $\mathcal{I}_{\alpha,l}$, we will condition on the event that $\|\mathbf{u}_i^\beta\|_\infty \leq f_j$ fails. We called this failure event $\mathcal{L}_{\alpha,l}$. The union of the failure events gives for α, l satisfying $l \notin \alpha$,

$$(10.91) \quad \begin{aligned} \mathcal{F}_{\alpha,l} &\subset (\mathcal{K}_{\alpha,l}) \cup (\mathcal{I}_{\alpha,l}) \\ &\subset (\mathcal{B}_{\alpha,l}) \cup (\mathcal{I}_{\alpha,l} \cap \mathcal{L}_{\alpha,l}) \cup (\mathcal{I}_{\alpha,l} \cap \overline{\mathcal{L}_{\alpha,l}}). \end{aligned}$$

Obviously, the intersection of two sets is contained in both sets, so $\mathcal{I}_{\alpha,l} \cap \mathcal{L}_{\alpha,l} \subset \mathcal{L}_{\alpha,l}$. By taking the union of both sides over $|\alpha| = j-1, l \notin \alpha$,

$$\bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} \mathcal{F}_{\alpha,l} \subset \bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} (\mathcal{B}_{\alpha,l}) \cup \bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} (\mathcal{L}_{\alpha,l}) \cup \bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} (\mathcal{I}_{\alpha,l} \cap \overline{\mathcal{L}_{\alpha,l}}).$$

For a given α , we have only considered the coordinates l such that $l \notin \alpha$. In order to bound γ_{j-1} , we must also consider the other coordinates (the ones included in α). This case is easy to handle, as by Lemma 10.3,

$$(10.92) \quad \bigcup_{\substack{|\alpha|=j-1 \\ l \in \alpha}} \mathcal{F}_{\alpha,l} \subset \bigcup_{|\alpha|=j-1} \mathcal{B}_{\alpha,E}.$$

By definition of the events $\mathcal{B}_j$,

$$\mathcal{B}_{j-1} = \bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} (\mathcal{B}_{\alpha,l}) \cup \bigcup_{|\alpha|=j-1} \mathcal{B}_{\alpha,E}.$$

Therefore,

$$\bigcup_{\substack{|\alpha|=j-1 \\ 1 \leq l \leq n}} \mathcal{F}_{\alpha,l} \subset \mathcal{B}_{j-1} \cup \bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} (\mathcal{L}_{\alpha,l}) \cup \bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} (\mathcal{I}_{\alpha,l} \cap \overline{\mathcal{L}_{\alpha,l}}).$$

By the definition of $\mathcal{I}_j$,

$$(10.93) \quad \bigcup_{\substack{|\alpha|=j-1 \\ 1 \leq l \leq n}} \mathcal{F}_{\alpha,l} \subset \mathcal{B}_{j-1} \cup \mathcal{I}_{j-1} \cup \bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} \mathcal{L}_{\alpha,l}.$$

We make two observations about (10.93). First, the event on the LHS of the inclusion is the event

$$\bigcup_{\substack{|\alpha|=j-1 \\ 1 \leq l \leq n}} \mathcal{F}_{\alpha,l} = \bigcup_{\substack{|\alpha|=j-1 \\ 1 \leq l \leq n}} \{|u_{il}^\alpha| > f_{j-1}\} = \left\{ \max_{\alpha, |\alpha|=j-1} \|\mathbf{u}_i^\alpha\|_\infty > f_{j-1} \right\}.$$

Second, because we are considering $l \notin \alpha$, the union of the $\mathcal{L}_{\alpha,l}$ on the RHS is precisely

$$\bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} \mathcal{L}_{\alpha,l} = \bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} \left\{ \left\| \mathbf{u}_i^{\alpha \cup \{l\}} \right\|_\infty > f_j \right\} = \left\{ \max_{\alpha, |\alpha|=j} \|\mathbf{u}_i^\alpha\|_\infty > f_j \right\}.$$

Therefore, by the definition of the γ_j ,

$$(10.94) \quad \mathbb{P}\left(\bigcup_{\substack{|\alpha|=j-1 \\ 1 \leq l \leq n}} \mathcal{F}_{\alpha,l}\right) = \gamma_{j-1}, \text{ and } \mathbb{P}\left(\bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} \mathcal{L}_{\alpha,l}\right) = \gamma_j.$$

By the union bound, (10.93) implies

$$(10.95) \quad \gamma_{j-1} \leq \mathbb{P}(\mathcal{B}_{j-1}) + \mathbb{P}(\mathcal{I}_{j-1}) + \gamma_j.$$

By Lemma 10.2,

$$(10.96) \quad \mathbb{P}(\mathcal{B}_{j-1}) \leq 65 \times 9^{2r} n^j \exp(-c_2 \log^2 n) + 2\tau.$$

By Lemma 10.5,

$$(10.97) \quad \mathbb{P}(\mathcal{I}_{j-1}) \leq 2n^j \exp(-c_2 \log^2 n).$$

The inequalities (10.95), (10.96), and (10.97) imply the lemma, since

$$\mathbb{P}(\mathcal{I}_{j-1}) + \mathbb{P}(\mathcal{B}_{j-1}) \leq 66 \times 9^{2r} n^j \exp(-c_2 \log^2 n) + 2\tau.$$

□

In the next two sections, we prove Lemmas 10.2, 10.3, 10.4, and 10.5.

11 Proof of Lemmas 10.3 and 10.4

To prove Lemma 10.3, we begin with a proposition, stated for deterministic H .

PROPOSITION 11.1. *Suppose H^α is a matrix equal to H , but whose rows and columns indexed by the index set α are set to zero. Suppose $l \in \alpha$, and that $\sigma_i > 2 \|H^\alpha\|$. Let $\mathbf{u}_i^\alpha$ be the i th singular vector of $A^\alpha = A + H^\alpha$. Then*

$$(11.98) \quad |u_{il}^\alpha| \leq 2\sqrt{r}\kappa_i \|U\|_\infty.$$

Proof of Lemma 10.3 given the proposition. Recall that we wish to show that

$$(11.99) \quad \bigcup_{\substack{|\alpha|=j \\ l \in \alpha}} \mathcal{F}_{\alpha,l} \subset \bigcup_{|\alpha|=j} \mathcal{B}_{\alpha,E}.$$

Consider an index set α such that $|\alpha| = j$, and an index l such that $l \in \alpha$. We will show that $\overline{\mathcal{B}_{\alpha,E}} \subset \overline{\mathcal{F}_{\alpha,l}}$. Suppose $\overline{\mathcal{B}_{\alpha,E}}$ holds, which means $\|E^\alpha\| \leq T$. We need to show that this implies $|u_{il}^\alpha| \leq f_j$, which means that $\overline{\mathcal{F}_{\alpha,l}}$ holds. By the $(c, \tau, 2)$ strong stability assumption (see (a) in Definition 3.1), $\sigma_i > cT$, and c is much larger than 2. Therefore,

$$\sigma_i \geq cT \geq c\|E^\alpha\| > 2\|E^\alpha\|.$$

Since $l \in \alpha$, the conditions of Proposition 11.1 are satisfied with the index set α and $H = E$. Proposition 11.1 implies that

$$(11.100) \quad |u_{il}^\alpha| \leq 2\sqrt{r}\kappa_i \|U\|_\infty \leq C_0\kappa_i \|U\|_\infty \leq f_j,$$

since $C_0 > 2\sqrt{r}$. $\square$

Proof of Proposition 11.1. Recall that u_{il}^α is the l th entry of a singular vector $\mathbf{u}_i^\alpha$ of A^α , corresponding to a singular value σ_i^α . By definition, we have

$$|u_{il}^\alpha| = |\mathbf{e}_l^T \mathbf{u}_i^\alpha| = \frac{1}{\sigma_i^\alpha} |\mathbf{e}_l^T A^\alpha \mathbf{u}_i^\alpha|.$$

Since the l th row of H^α is zero, A^α and A have the same l th row, which implies that we can replace A^α by A to obtain

$$|u_{il}^\alpha| = \frac{1}{\sigma_i^\alpha} |\mathbf{e}_l^T A \mathbf{u}_i^\alpha|.$$

By Fact 5.1, Fact 5.2, and the assumption of the proposition, it is easy to deduce that $\sigma_i^\alpha > \sigma_i/2$. Writing, $A = U\Sigma U^T$ using the spectral decomposition, we obtain

$$|u_{il}^\alpha| \leq 2\sigma_i^{-1} |\mathbf{e}_l^T U \Sigma U^T \mathbf{u}_i^\alpha|.$$

Notice that

$$|\mathbf{e}_l^T U \Sigma U^T \mathbf{u}_i^\alpha| \leq \|\mathbf{e}_l^T U\|_2 \|\Sigma\| \|U^T \mathbf{u}_i^\alpha\|_2 \leq \sigma_1 \|U_{l,\cdot}\|_2,$$

This is because $\|\Sigma\| = \sigma_1$, $\mathbf{e}_l^T U = U_{l,\cdot}$, the l th row of U , and $\|U \mathbf{u}_i^\alpha\|_2 \leq 1$. This and the previous bound imply

$$|u_{il}^\alpha| \leq 2 \frac{\sigma_1}{\sigma_i} \|U_{l,\cdot}\|_2 = 2\kappa_i \|U_{l,\cdot}\|_2 \leq 2\kappa_i \sqrt{r} \|U\|_\infty.$$

where the last line uses that $\|U_{l,\cdot}\|_2 \leq \sqrt{r} \|U\|_\infty$ by the Cauchy-Schwarz inequality. $\square$

We move to the proof of Lemma 10.4. The following lemma is used to prove Lemma 10.4 by showing that the conditions for Theorem 3.2 hold on the complement of $\mathcal{B}_{\alpha,l}$. We use Theorem 3.2 to show that $\mathcal{K}_{\alpha,l} \subset \mathcal{B}_{\alpha,l}$.

LEMMA 11.1. (THEOREM 3.2 CAN BE APPLIED ON $\overline{\mathcal{B}_{\alpha,l}}$) Let $0 \leq j \leq j^*$. Let α be an index set satisfying $|\alpha| = j$ and let l be an index such that $l \notin \alpha$. Under the conditions of Lemma 3.1, if $\overline{\mathcal{B}_{\alpha,l}}$ occurs, the conditions for Theorem 3.2 hold with A , $H = E^\alpha$, and coordinate l .

Proof of Lemma 10.4 given Lemma 11.1 Let α be an index set such that $|\alpha| = j$, and let $l \notin \alpha$. Recall that we wish to show that $\mathcal{K}_{\alpha,l} \subset \mathcal{B}_{\alpha,l}$. Set $\beta = \alpha \cup \{l\}$. Suppose $\overline{\mathcal{B}_{\alpha,l}}$ holds. We wish to show that $\overline{\mathcal{K}_{\alpha,l}}$ holds, which is equivalent to showing that

$$(11.101) \quad |u_{il}^\alpha| \leq 3C_0\kappa_i\|U\|_\infty + 256r \frac{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle|}{\sigma_i}.$$

By Lemma 11.1 we can apply Theorem 3.2 on $\overline{\mathcal{B}_{\alpha,l}}$ with A and $H = E^\alpha$ for coordinate l . This gives

$$(11.102) \quad |u_{il}^\alpha - u_{il}| \leq C_0 \|U_{l,\cdot}\|_\infty \left[\kappa_i \|\mathbf{u}_i^\alpha - \mathbf{u}_i\|_2 + \epsilon_1^\alpha(i) + a_l^\alpha \kappa_i \epsilon_2(i) \right] + 256r \frac{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle|}{\sigma_i},$$

where $a_l^\alpha = \|U^{\beta T} \mathbf{x}(\alpha, l)\|_2$ and $\epsilon_1^\alpha(i) = \frac{\|E^\alpha\|}{\sigma_i}$. We will bound the term in the large brackets on the RHS and show that it is smaller than $\frac{5}{2}\kappa_i$. We will show that on $\overline{\mathcal{B}_{\alpha,l}}$,

$$(11.103) \quad \epsilon_1^\alpha(i) + a_l^\alpha \kappa_i \epsilon_2(i) \leq \frac{\kappa_i}{2}.$$

Assume for now that this holds. By using the trivial bound

$$(11.104) \quad \|\mathbf{u}_i^\alpha - \mathbf{u}_i\|_2 \leq 2,$$

(11.102) and (11.103) imply that

$$(11.105) \quad |u_{il}^\alpha - u_{il}| \leq \frac{5}{2}C_0\kappa_i\|U\|_\infty + 256r \frac{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle|}{\sigma_i}.$$

Moving $|u_{il}|$, which is smaller than $\|U\|_\infty$, to the right, we obtain

$$(11.106) \quad |u_{il}^\alpha| \leq \left(\frac{5}{2}C_0 + 1\right)\kappa_i\|U\|_\infty + 256r \frac{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle|}{\sigma_i}.$$

Since $C_0 > 2$, we obtain (11.101), proving the lemma.

Therefore, what remains is to verify (11.103) on $\overline{\mathcal{B}_{\alpha,l}}$. Assume that $\overline{\mathcal{B}_{\alpha,l}}$ holds. Then, by definition of $\mathcal{B}_{\alpha,E}$, we have $\|E^\alpha\| \leq T$, which implies that

$$(11.107) \quad \epsilon_1^\alpha(i) = \frac{\|E^\alpha\|}{\sigma_i} \leq \frac{T}{\sigma_i} < \frac{1}{4}.$$

The last inequality uses the $(c, \tau, 2)$ stability assumption, which gives $\sigma_i > cT > 4T$. For $a_l^\alpha \kappa_i \epsilon_2(i)$, recall that $\epsilon_2(i) = \frac{1}{\delta_i}$. Furthermore, on $\overline{\mathcal{B}_{\alpha,l}}$, the definition of $\mathcal{B}_{\alpha,l,2}$ guarantees that $a_l^\alpha \leq K\sqrt{2rc_2 \log^2 n}$. Therefore,

$$(11.108) \quad a_l^\alpha \epsilon_2(i) \leq \frac{K\sqrt{2rc_2 \log^2 n}}{\delta_i} < \frac{1}{4}.$$

The last inequality uses $(c, \tau, 2)$ stability, which guarantees $\delta_i > cK\sqrt{\log^2 n}$, and the fact that $c > 4\sqrt{2rc_2}$. We conclude that

$$(11.109) \quad \epsilon_1^\alpha(i) + a_l^\alpha \kappa_i \epsilon_2(i) \leq \frac{\kappa_i}{2}.$$

This verifies (11.103) and thus completes the proof. $\square$

Proof of Lemma 11.1 Let α be an index set α with $|\alpha| = j$, and let l be an index such that $l \notin \alpha$. Set $\beta = \alpha \cup \{l\}$. Suppose that $\overline{\mathcal{B}}_{\alpha,l}$ holds. Recall that we want to show that this implies that the conditions for Theorem 3.2 are satisfied with $H = E^\alpha$ and coordinate l .

Assume $\overline{\mathcal{B}}_{\alpha,l}$ occurs. We first have to check that this implies that

$$(11.110) \quad \min\{|\sigma_i^\alpha - \sigma_{i+1}^\beta|, |\sigma_i^\alpha - \sigma_{i-1}^\beta|\} \geq \delta_i/2.$$

It is clear that this is true by definition of $\mathcal{B}_{\alpha,l,1}$. Then, we have to check that $\overline{\mathcal{B}}_{\alpha,l}$ implies

$$(11.111) \quad \sigma_i > C_0 \|E^\alpha\|.$$

On $\overline{\mathcal{B}}_{\alpha,l}$, $\|E^\alpha\| \leq T$ (see $\mathcal{B}_{\alpha,E}$). By $(c, \tau, 2)$ stability, $\sigma_i > cT$. Therefore,

$$\sigma_i > cT > C_0 \|E^\alpha\|,$$

as desired. Finally, we verify that

$$(11.112) \quad \delta_i > C_0 \max \left\{ \|U^{\beta T} \mathbf{x}(\alpha, l)\|_2, \kappa_i \|E^\alpha\| \|U\|_\infty \right\}.$$

On $\overline{\mathcal{B}}_{\alpha,l}$ (see $\mathcal{B}_{\alpha,l,2}$),

$$\|U^{\beta T} \mathbf{x}(\alpha, l)\|_2 \leq K \sqrt{2rc_2 \log^2 n}.$$

Because of the $(c, \tau, 2)$ stability assumption,

$$\delta_i > c \max \{K \sqrt{\log^2 n}, T \kappa_i \|U\|_\infty\},$$

which ensures that (11.112) holds, because $c > C_0 \sqrt{2rc_2}$. $\square$

12 Proof of Lemmas 10.2 and 10.5

Proof of Lemma 10.2 Let $0 \leq j \leq j^*$. Recall that

$$\mathcal{B}_j := \bigcup_{\substack{|\alpha|=j \\ l \notin \alpha}} (\mathcal{B}_{\alpha,l,1} \cup \mathcal{B}_{\alpha,l,2}) \cup \bigcup_{|\alpha|=j} \mathcal{B}_{\alpha,E}.$$

We will continue to use $\beta = \beta(\alpha, l) = \alpha \cup \{l\}$. In particular, β depends on both α and l . We observe that $\mathcal{B}_j \subset \mathcal{B}_{j1} \cup \mathcal{B}_{j2} \cup \mathcal{B}_{jE}$, where

$$(12.113) \quad \begin{aligned} \mathcal{B}_{j1} &:= \bigcup_{|\alpha|=j} \bigcup_{l \notin \alpha} \{ \min\{|\sigma_i^\alpha - \sigma_{i+1}^\beta|, |\sigma_i^\alpha - \sigma_{i-1}^\beta|\} < \delta_i/2 \}, \\ \mathcal{B}_{j2} &:= \bigcup_{|\alpha|=j} \bigcup_{l \notin \alpha} \left\{ \|U^{\beta T} \mathbf{x}(\alpha, l)\|_2 \geq K \sqrt{2c_2 r \log^2 n} \right\}, \text{ and} \\ \mathcal{B}_{jE} &:= \bigcup_{|\alpha|=j} \{ \|E^\alpha\| > T \}. \end{aligned}$$

We will bound the probabilities of these three events and use the union bound to conclude.

Probability of $\mathcal{B}_{jE}$. By Fact 5.2 for all $|\alpha| = j$, we have the deterministic bound $\|E^\alpha\| \leq \|E\|$. Therefore, defining

$$\mathcal{B}_E := \{\|E\| \geq T\}, \text{ we have}$$

$$(12.114) \quad \mathbb{P}(\mathcal{B}_{jE}) \leq \mathbb{P}(\mathcal{B}_E) \leq \tau,$$

by definition of T .

Probability of $\mathcal{B}_{j2}$. Recall that

$$\mathcal{B}_{j2} = \cup_{|\alpha|=j} \cup_{l \notin \alpha} \left\{ \|U^{\beta T} \mathbf{x}(\alpha, l)\|_2 \geq K \sqrt{2c_2 r \log^2 n} \right\}.$$

Let α be an index set such that $|\alpha| = j$ and let $l \notin \alpha$. Set $\beta = \alpha \cup \{l\}$.

We use Hoeffding's inequality. Observe that $\mathbf{x}(\alpha, l)$ and A^β are independent. Therefore, each entry of the length r vector $U^{\beta T} \mathbf{x}(\alpha, l)$ is the inner product of a K bounded random vector with independent entries (the vector $\mathbf{x}(\alpha, l)$) with a unit vector from which it is independent. By applying Corollary 5.1 to each entry, for $1 \leq k \leq r$,

$$(12.115) \quad \mathbb{P}\{|[U^{\beta T} \mathbf{x}(\alpha, l)]_k| \geq K \sqrt{2c_2 \log^2 n}\} \leq 2 \exp(-c_2 \log^2 n).$$

Therefore, by taking the union bound over $1 \leq k \leq r$,

$$(12.116) \quad \mathbb{P}\left\{ \|U^{\beta T} \mathbf{x}(\alpha, l)\|_2 \geq K \sqrt{2rc_2 \log^2 n} \right\} \leq 2r \exp(-c_2 \log^2 n).$$

By taking the union bound of (12.116) over $|\alpha| = j$, and $l \notin \alpha$,

$$(12.117) \quad \mathbb{P}[\mathcal{B}_{j2}] \leq 2rn^{j+1} \exp(-c_2 \log^2 n).$$

Probability of $\mathcal{B}_{j1}$. Recall that

$$\mathcal{B}_{j1} = \cup_{|\alpha|=j} \cup_{l \notin \alpha} \left\{ \min\{|\sigma_i^\alpha - \sigma_{i+1}^\beta|, |\sigma_i^\alpha - \sigma_{i-1}^\beta|\} < \delta_i/2 \right\}.$$

For an index set α such that $|\alpha| = j$ and a coordinate $l \notin \alpha$, set $\beta = \alpha \cup \{l\}$. Similar to the proof of Theorem 3.3 (with σ_i^α playing the role of $\tilde{\sigma}_i$ and σ_i^β playing the role of $\sigma_i^{\{l\}}$), define the events $\mathcal{G}_{j,i+1}$ and $\mathcal{G}_{j,i-1}$

$$(12.118) \quad \begin{aligned} \mathcal{G}_{j,i+1} &:= \bigcap_{|\alpha|=j} \left\{ \max_{k=i,i+1} \max_{l \notin \alpha} \{|\sigma_k^\beta - \sigma_k|, |\sigma_k^\alpha - \sigma_k|\} \leq 24r \left[K \sqrt{rc_2 \log^2 n} + \frac{\|E\|^2}{\tilde{\sigma}_k} + \frac{\|E\|^3}{\tilde{\sigma}_k^2} \right] \right\}, \\ \mathcal{G}_{j,i-1} &:= \bigcap_{|\alpha|=j} \left\{ \max_{k=i,i-1} \max_{l \notin \alpha} \{|\sigma_k^\beta - \sigma_k|, |\sigma_k^\alpha - \sigma_k|\} \leq 24r \left[K \sqrt{rc_2 \log^2 n} + \frac{\|E\|^2}{\tilde{\sigma}_k} + \frac{\|E\|^3}{\tilde{\sigma}_k^2} \right] \right\}. \end{aligned}$$

These events keep the relevant singular values of the perturbations A^α and A^β for all α, l close to the original singular values to ensure the gap remains large. Recall that we applied Theorem 3.2 for $H = E^\alpha$ and $H^{\{l\}} = E^\beta$. The singular values of both A^α and A^β are controlled by $\mathcal{G}_{j,i-1}$ and $\mathcal{G}_{j,i+1}$.

Using the same argument we employed in the proof of Proposition 9.1, it is straightforward to show using case analysis and $(c, \tau, 2)$ stability that

$$(12.119) \quad \mathbb{P}(\mathcal{B}_{j1}) \leq \mathbb{P}(\overline{\mathcal{G}_{j,i-1}}) + \mathbb{P}(\overline{\mathcal{G}_{j,i+1}}) + \mathbb{P}(\mathcal{B}_E),$$

where we recall $\mathcal{B}_E = \{\|E\| \geq T\}$. $\mathcal{B}_E$ is an event which has probability at most τ by definition of T . Bounding the probabilities of $\overline{\mathcal{G}_{j,i-1}}$ and $\overline{\mathcal{G}_{j,i+1}}$ proceeds virtually identically to the proof of Lemma 9.2, so we omit the details of the calculation. The probabilities of $\overline{\mathcal{G}_{j,i-1}}$ and $\overline{\mathcal{G}_{j,i+1}}$ will be bounded using the result of 5.1, which we recall in Theorem 5.1, applied to the random perturbations E^α and E^β . Both have norm at most $\|E\|$ by Fact 5.2. By Theorem 5.1 applied with $t = K \sqrt{128rc_2 \log^2 n}$ to E^α and E^β , and the union bound over $|\alpha| = j$, $l \notin \alpha$,

$$(12.120) \quad \mathbb{P}(\overline{\mathcal{G}_{j,i-1}}) + \mathbb{P}(\overline{\mathcal{G}_{j,i+1}}) \leq 64 \times 9^{2r} n^{j+1} \exp(-c_2 \log^2 n).$$

Together with (12.119), this implies

$$(12.121) \quad \mathbb{P}(\mathcal{B}_{j1}) \leq 64 \times 9^{2r} n^{j+1} \exp(-c_2 \log^2 n) + \tau.$$

We have now bounded $\mathcal{B}_{j1}$, $\mathcal{B}_{j2}$, and $\mathcal{B}_{jE}$. To conclude the proof, use the union bound and (12.114), (12.117), and (12.121). This gives

$$\mathbb{P}(\mathcal{B}_j) \leq (64 \times 9^{2r} + 2r)n^{j+1} \exp(-c_2 \log^2 n) + \tau \leq 65 \times 9^{2r} n^{j+1} \exp(-c_2 \log^2 n) + \tau.$$

□

What remains is to prove Lemma 10.5, which bounds the probability of $\mathcal{I}_j$. To begin, we reproduce the definitions of the relevant events. Let α be an index set with $|\alpha| = j$ and let $1 \leq l \leq n$. Recall that $\beta = \alpha \cup l$, and we defined the events

$$\mathcal{L}_{\alpha,l} = \left\{ \left\| \mathbf{u}_i^\beta \right\|_\infty > f_{j+1} \right\},$$

$$\mathcal{I}_{\alpha,l} = \{ \langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle \geq c_2 \sqrt{2Kn} f_{j+1} \log^2 n \}, \text{ and}$$

$$\mathcal{I}_j = \bigcup_{\substack{|\alpha|=j \\ l \notin \alpha}} (\mathcal{I}_{\alpha,l} \cap \overline{\mathcal{L}_{\alpha,l}}).$$

We will need the following lemma. We first introduce some notation. The notation $Y \in \overline{\mathcal{L}_{\alpha,l}}$ means that Y is a possible realization of E^β such that $\overline{\mathcal{L}_{\alpha,l}}$ holds.

LEMMA 12.1. *Let α be an index set with $|\alpha| = j$, and let l be a coordinate such that $l \notin \alpha$. Set $\beta = \alpha \cup \{l\}$. Let $Y \in \overline{\mathcal{L}_{\alpha,l}}$. Then,*

$$(12.122) \quad \mathbb{P}(\mathcal{I}_{\alpha,l} | E^\beta = Y) \leq 2 \exp(-c_2 \log^2 n).$$

Proof of Lemma 10.5 given Lemma 12.1. The goal is to bound the the probability of

$$\mathcal{I}_j = \bigcup_{\substack{|\alpha|=j-1 \\ l \notin \alpha}} (\mathcal{I}_{\alpha,l} \cap \overline{\mathcal{L}_{\alpha,l}}).$$

By conditioning on $\overline{\mathcal{L}_{\alpha,l}}$,

$$(12.123) \quad \mathbb{P}(\mathcal{I}_{\alpha,l} \cap \overline{\mathcal{L}_{\alpha,l}}) = \mathbb{P}(\mathcal{I}_{\alpha,l} | \overline{\mathcal{L}_{\alpha,l}}) \mathbb{P}(\overline{\mathcal{L}_{\alpha,l}}) \leq \mathbb{P}(\mathcal{I}_{\alpha,l} | \overline{\mathcal{L}_{\alpha,l}}),$$

where we use the trivial bound $\mathbb{P}(\overline{\mathcal{L}_{\alpha,l}}) \leq 1$. We now bound the RHS to obtain

$$(12.124) \quad \mathbb{P}(\mathcal{I}_{\alpha,l} \cap \overline{\mathcal{L}_{\alpha,l}}) \leq \mathbb{P}(\mathcal{I}_{\alpha,l} | \overline{\mathcal{L}_{\alpha,l}}) \leq \sup_{Y \in \overline{\mathcal{L}_{\alpha,l}}} \mathbb{P}(\mathcal{I}_{\alpha,l} | E^\beta = Y).$$

Lemma 12.1 thus implies that

$$\mathbb{P}(\mathcal{I}_{\alpha,l} \cap \overline{\mathcal{L}_{\alpha,l}}) \leq 2 \exp(-c_2 \log^2 n).$$

The union bound over $|\alpha| = j$ and $l \notin \alpha$ gives

$$(12.125) \quad \mathbb{P}(\mathcal{I}_j) \leq 2n^{j+1} \exp(-c_2 \log^2 n).$$

□

Proof of Lemma 12.1. Recall that we are considering index set α satisfying $|\alpha| = j$ and a coordinate $l \notin \alpha$. Set $\beta = \alpha \cup \{l\}$. We are looking for a bound of

$$\mathbb{P}(\mathcal{I}_{\alpha,l}|E^\beta = Y) = \mathbb{P}\left\{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle| \geq c_2 \sqrt{2Kn} f_{j+1} \log^2 n \mid E^\beta = Y\right\},$$

where $Y \in \overline{\mathcal{L}_{\alpha,l}}$. In other words, Y is a possible realization of E^β satisfying $\|\mathbf{u}_i^\beta\|_\infty \leq f_{j+1}$. Recall that $\mathbf{x}(\alpha, l)$ is the l th row of E^α (with l th entry divided by 2). Conditional on E^β equalling such a Y , $\mathbf{u}_i^\beta$ is a deterministic unit vector whose entries have absolute value at most f_{j+1} . The only randomness in each event thus comes from the vector $\mathbf{x}(\alpha, l)$. Let us name the entries of $\mathbf{x}(\alpha, l)$ as x_k . It follows that the inner product $\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle$, conditional on $E^\beta = Y$, is the sum of independent, Kf_{j+1} bounded, mean zero random variables $x_k u_{ik}^\beta$. Thus, this quantity can be bounded with Bernstein's inequality (Lemma 5.2).

We are applying Bernstein's inequality conditionally, so we also need to find a bound for the sum of the conditional second moments of the $x_k u_{ik}^\beta$. Since $\mathbf{u}_i^\beta$ is deterministic when we condition on E^β , and $\mathbf{x}(\alpha, l)$ is independent of E^β , we have

$$(12.126) \quad \sum_{k=1}^n \mathbb{E}\left[x_k^2 u_{ik}^{\beta 2} \mid E^\beta = Y\right] = \sum_{k=1}^n u_{ik}^{\beta 2} \mathbb{E}[x_k^2] \leq K \sum_{k=1}^n u_{ik}^{\beta 2} = K.$$

For the inequality, we use that the second moments of the entries of E are at most K by Assumption 3.1. The last equality uses the fact that $\mathbf{u}_i^\beta$ is a unit vector. Applying Bernstein's inequality then gives that

$$(12.127) \quad \begin{aligned} \mathbb{P}\left\{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle| > t \mid E^\beta = Y\right\} &\leq 2 \exp\left(\frac{-t^2/2}{\sum_{k=1}^n \mathbb{E}\left[x_k^2 u_{ik}^{\beta 2} \mid E^\beta = Y\right] + Kf_{j+1}t/3}\right) \\ &\leq 2 \exp\left(\frac{-t^2/2}{K + Kf_{j+1}t/3}\right). \end{aligned}$$

Setting $t = c_2 \sqrt{2Kn} f_{j+1} \log^2 n$,

$$(12.128) \quad \mathbb{P}\left\{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle| > c_2 \sqrt{2Kn} f_{j+1} \log^2 n \mid E^\beta = Y\right\} \leq 2 \exp\left(\frac{-c_2^2 Kn f_{j+1}^2 \log^4 n}{K + \frac{\sqrt{2}c_2}{3} K f_{j+1}^2 \sqrt{Kn} \log^2 n}\right).$$

We now upper bound the terms in the denominator. For the first term, since $n^{-1/2} \leq \|U\|_\infty \leq f_{j+1}$, we obtain that $K \leq Kn f_{j+1}^2$. For the second term, we use that $K \leq n$ by Assumption 3.1 so $\sqrt{Kn} \leq n$. This gives

$$(12.129) \quad \mathbb{P}\left\{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle| > c_2 \sqrt{2Kn} f_{j+1} \log^2 n \mid E^\beta = Y\right\} \leq 2 \exp\left(\frac{-c_2^2 Kn f_{j+1}^2 \log^4 n}{Kn f_{j+1}^2 + \frac{\sqrt{2}c_2}{3} K f_{j+1}^2 n \log^2 n}\right).$$

Since $c_2 = c_0 + 1 > 1$, we can upper bound the first term in the denominator on the RHS with $Kn f_{j+1}^2 \leq \frac{c_2}{3} Kn f_{j+1}^2 \log^2 n$. This gives

$$(12.130) \quad \begin{aligned} \mathbb{P}\left\{|\langle \mathbf{u}_i^\beta, \mathbf{x}(\alpha, l) \rangle| > c_2 \sqrt{2Kn} f_{j+1} \log^2 n \mid E^\beta = Y\right\} &\leq 2 \exp\left(\frac{-c_2^2 Kn f_{j+1}^2 \log^4 n}{(\frac{1+\sqrt{2}}{3})c_2 Kn f_{j+1}^2 \log^2 n}\right) \\ &\leq 2 \exp(-c_2 \log^2 n). \end{aligned}$$

□

13 The proof of Theorem 3.4

The proof follows fairly easily from the proof of Lemma 3.1. We only need to augment the events to reduce the wasteful $\log^2 n$ terms in the bounds for inner products. Recall the quantities $U^{\{l\}}$, $\mathbf{u}_i^{\{l\}}$, and $\mathbf{x}(l)$ from the proof of Theorem 3.3. Define the events

$$\mathcal{I}'_l := \{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x}(l) \rangle| \geq c_2 K \sqrt{2Kn} f_1 \log n\}.$$

Let

$$\mathcal{I}' := \bigcup_{1 \leq l \leq n} \mathcal{I}'_l.$$

Recall that in the proof of Lemma 3.1 we defined the events $\mathcal{B}_j$. We will be considering $\mathcal{B}_0$. Lastly, define

$$\mathcal{B}' := \bigcup_{1 \leq l \leq n} \left\{ \left\| U^{\{l\}T} \mathbf{x}(l) \right\|_2 \geq K \sqrt{2rc_2 \log n} \right\}.$$

We slightly abuse notation, as $\mathcal{B}'$ was defined and bounded in the proof of Theorem 3.3 (under the name $\mathcal{B}_2$). Define the failure event

$$\mathcal{F} = \mathcal{I}' \cup \mathcal{B}' \cup \mathcal{B}_0.$$

LEMMA 13.1.

$$(13.131) \quad \mathbb{P}(\mathcal{F}) \leq \tau \log n + 150 \times 9^{2r} n^{-c_0}.$$

Proof of Theorem 3.4 given Lemma 13.1. Let $1 \leq l \leq n$. By Lemma 13.1, $\mathcal{F}$ has probability at most $\tau \log n + 150 \times n^{-c_0}$. By Lemma 11.1, the conditions for Theorem 3.2 hold on $\bar{F} \subset \bar{B}_0$ with $A, H = E$, for coordinate l . The theorem gives the bound

$$(13.132) \quad |\tilde{u}_{il} - u_{il}| \leq C_0 \|U_{l,\cdot}\|_\infty \left[\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + a_l \kappa_i \epsilon_2(i) \right] + 256r \frac{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x}(l) \rangle|}{\sigma_i},$$

where we recall that $a_l = \|U^{\{l\}T} \mathbf{x}(l)\|_2$. On $\bar{F}$, we have $a_l \leq K \sqrt{2rc_2 \log n}$ because of $\mathcal{B}'$. Further, we have, by definition of $\mathcal{I}'$,

$$(13.133) \quad |\langle \mathbf{u}_i^{\{l\}}, \mathbf{x}(l) \rangle| \leq c_2 K \sqrt{2Kn} f_1 \log n \leq 4C_0 c_2 \sqrt{2Kn} \kappa_i \|U\|_\infty \log n,$$

where we use our previous observation that $f_1 \leq 4C_0 \kappa_i \|U\|_\infty$. By definition of c , we therefore have

$$|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x}(l) \rangle| < \frac{c}{256r} \kappa_i \|U\|_\infty \sqrt{Kn} \log n.$$

It follows that

$$(13.134) \quad |\tilde{u}_{il} - u_{il}| \leq c \|U_{l,\cdot}\|_\infty \left[\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + \kappa_i \epsilon_2(i) K \sqrt{\log n} \right] + \frac{c \kappa_i \|U\|_\infty \sqrt{Kn} \log n}{\sigma_i}.$$

Since l was arbitrary, we have

$$(13.135) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty \leq c \|U\|_\infty \left[\kappa_i \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \epsilon_1(i) + \kappa_i \epsilon_2(i) K \sqrt{\log n} \right] + \frac{c \kappa_i \|U\|_\infty \sqrt{Kn} \log n}{\sigma_i}.$$

□

13.1 Proof of Lemma 13.1. We will make use of the fact that $c_2 = (1 + c_0)$, so $\exp(-c_2 \log n) = n^{-c_0-1}$, for example. We first bound $\mathbb{P}(\mathcal{B}_0)$ and $\mathbb{P}(\mathcal{B}')$. By Lemma 10.2

$$\mathbb{P}(\mathcal{B}_j) \leq 65 \times 9^{2r} n^{j+1} \exp(-c_2 \log^2 n) + 2\tau,$$

so we have

$$\begin{aligned} \mathbb{P}(\mathcal{B}_0) &\leq 65 \times 9^{2r} n \exp(-c_2 \log^2 n) + 2\tau \\ &\leq 65 \times 9^{2r} n^{-c_2 \log n} + 2\tau \\ &\leq 65 \times 9^{2r} n^{-\log n} n^{-c_0} + 2\tau \\ &\leq 9^{2r} n^{-c_0} + 2\tau. \end{aligned} \tag{13.136}$$

As we mentioned, in the poof of Theorem 3.3, we bounded (see 9.63)

$$\mathbb{P}(\mathcal{B}') \leq 2rn^{-c_0}. \tag{13.137}$$

We will bound $\mathbb{P}(\mathcal{I}')$ in the proceeding lemma. Recall that $j^* = \lceil \frac{50 \log n}{\log(\log n)} \rceil + 3$.

LEMMA 13.2.

$$\mathbb{P}(\mathcal{I}') \leq 2j^* \tau + 133 \times 9^{2r} n^{-c_0}. \tag{13.138}$$

The union bound, Lemma 13.2, (13.136), and (13.137) show that

$$\mathbb{P}(\mathcal{F}) = \mathbb{P}(\mathcal{I}' \cup \mathcal{B}_0 \cup \mathcal{B}') \leq (134 \times 9^{2r} + 2r)n^{-c_0} + (2j^* + 2)\tau.$$

Since $j^* = o(\log n)$, we have

$$\mathbb{P}(\mathcal{F}) \leq 150 \times 9^{2r} n^{-c_0} + \tau \log n,$$

as desired (with room to spare). To complete the proof of Lemma 13.1, we must prove Lemma 13.2

Before proceeding with the proof, we recall some notation. For an index set α of size j and a coordinate $1 \leq l \leq n$, we defined the event $\mathcal{L}_{\alpha,l} = \{\|\mathbf{u}_i^{\alpha \cup \{l\}}\|_\infty > f_{j+1}\}$ in the proof of Lemma 10.1. We will consider these events with $\alpha = \{\}$ in the following proposition.

PROPOSITION 13.1. *Let*

$$\mathcal{J} = \bigcup_{1 \leq l \leq n} (\mathcal{I}'_l \cap \overline{\mathcal{L}_{\{\},l}}).$$

Then,

$$\mathbb{P}(\mathcal{J}) \leq 2n \exp(-c_2 \log n). \tag{13.139}$$

The proof of Proposition 13.1 is a repetition of the computations in Lemma 10.5 using Bernstein's inequality conditionally. We place the details in Appendix C

Proof of Lemma 13.2. We complete the task of bounding $\mathbb{P}(\mathcal{I}')$. By conditioning on $\mathcal{L}_{\{\},l}$, we have

$$\mathcal{I}'_l \subset \mathcal{L}_{\{\},l} \cup (\mathcal{I}'_l \cap \overline{\mathcal{L}_{\{\},l}}). \tag{13.140}$$

Therefore,

$$\mathcal{I}' \subset \bigcup_{1 \leq l \leq n} \mathcal{L}_{\{\},l} \cup \mathcal{J}. \tag{13.141}$$

The probability of $\mathcal{J}$ is bounded using the Proposition 13.1.

$$\mathbb{P}(\mathcal{J}) \leq 2n \exp(-c_2 \log n) = 2n^{-c_0}. \tag{13.142}$$

Next, we recognize the other event on the RHS of (13.141).

$$(13.143) \quad \mathbb{P}\left(\bigcup_{1 \leq l \leq n} \mathcal{L}_{\{l\},l}\right) = \mathbb{P}\{\max_l \|\mathbf{u}^{\{l\}}\|_\infty > f_1\} = \gamma_1.$$

The last equality is the definition of γ_1 . We bounded the γ_j iteratively in Lemma 10.1. A routine calculation virtually identical to the one done in (10.81) gives the following bound for γ_1 .

$$(13.144) \quad \mathbb{P}\left(\bigcup_{1 \leq l \leq n} \mathcal{L}_{\{l\},l}\right) = \gamma_1 \leq 2j^*\tau + 132 \times 9^{2r}n^{-c_0}.$$

The bounds (13.142) and (13.144) show that

$$\mathbb{P}(\mathcal{I}') \leq (132 \times 9^{2r} + 2)n^{-c_0} + 2j^*\tau \leq 133 \times 9^{2r}n^{-c_0} + 2j^*\tau,$$

as desired.

□

14 Application: A simple algorithm for clustering problems

A number of clustering problems have the following common form. A vertex set V is partitioned into r subsets $V_1, \dots, V_r$, and between each pair V_i, V_j we draw edges independently with probability p_{ij} (we allow $i = j$). The task is to find a particular set V_j or all the parts $V_1, \dots, V_r$ given one instance of the random graph [7, 13, 28, 34, 38, 39, 40, 41, 50].

The most popular approach to this problem is spectral, which typically consists of two steps. In the first step, one considers the coordinates of a singular vector of the adjacency matrix of the graph (or more generally the projection of the row vectors of the adjacency matrix onto a low dimensional singular space), and run a standard clustering algorithm. The output of this step is an approximation of the truth. In the second step, one applies adhoc combinatorial techniques to clean up the output, in order to recover the mis-classified vertices.

It has been conjectured that in many cases, the cleaning step is not necessary. Our result makes an important contribution towards solving this problem. The critical point here is that the existence of misclassified vertices, in many settings, is just an artifact of the analysis in the first step, which typically relies on ℓ_2 norm estimates. Notice that any ℓ_2 norm estimate, even sharp, could only imply that a majority of the vertices are well classified, and this leads to the necessity of the second step. Once we have a strong ℓ_∞ norm estimate, then we would be able to classify all the vertices at once.

As we stated in Section 4, our new infinity norm estimates enable us to overcome the shortcomings in clustering algorithms that rely on ℓ_2 analysis in a number of settings. This results in fast and simple new algorithms for a wide variety of problems. All matrices in this section are positive semi-definite, so there is no difference between singular vectors and eigenvectors.

14.1 The hidden clique problem. The (simplest form) of the hidden clique problem is the following: Hide a clique X of size k in the random graph $G(n, 1/2)$. Can we find X in polynomial time?

Notice that the largest clique in $G(n, 1/2)$, with overwhelming probability, has size approximately $2 \log n$ [8]. Thus, for any k bigger than $(2 + \epsilon) \log n$, with any constant $\epsilon > 0$, X would be abnormally large and therefore detectable, by brute-force at least. For instance, one can check all vertex sets of size k to see if any of them form a clique. However, finding X in polynomial time is a different matter, and the best current bound for k is $k \geq c\sqrt{n}$, for any constant $c > 0$. This was first achieved by Alon, Krivelevich, and Sudakov [7]; see also [40, 30] for later developments concerning faster algorithms for certain values of c .

The Alon-Krivelevich-Sudakov algorithm runs as follows. It first finds X when c is sufficiently large, then uses a simple sampling trick to reduce the case of small c to this case.

To find the clique for a large c , they first compute the second eigenvector of the adjacency matrix of the graph and locate the first largest k coordinates in absolute value. Call this set Y . This is an approximation of the clique X , but not yet totally accurate. The second, combinatorial, step is to define the set X as the vertices

in the graph with at least $3/4k$ neighbors in Y . The authors then proved that with high probability, X is indeed the hidden clique.

With our new results, we can find X immediately by a slightly modified version of the first step, omitting the second combinatorial step. Before starting the main step of the algorithm, we change all zeros in the adjacency matrix to -1 .

ALGORITHM 14.1. (FIRST SINGULAR VECTOR CLUSTERING-FSC) Compute the first singular vector. Let x be the largest value of the coordinates and let X be the set of all coordinates with value at least $x/2$.

This is perhaps the simplest algorithm for this problem, as computing the first singular vector of a large matrix is a routine operation that appears in all standard numerical linear algebra packages.

THEOREM 14.1. *There is a constant c_0 such that for all $k \geq c_0\sqrt{n}$, FSC outputs the hidden clique correctly with probability at least .99.*

Proof of Theorem 14.1. After the switching of zeroes to minus ones, the adjacency matrix $\tilde{A}$ has an all one block of size k (corresponding the hidden cliques), and the rest are ± 1 bits. For convenience, we assume that the all-one block is at the left-top corner. Thus, we can write $\tilde{A} = A + E$, where A has an all-one block on its leading principal sub-matrix of size k and the rest of the entries are zero. E is a random matrix with ± 1 entries with a zero block of size k .

Notice that the matrix A has rank 1, with $\sigma_1 = k$, and first singular vector

$$(1/\sqrt{k}, \dots, 1/\sqrt{k}, 0, 0, \dots, 0).$$

So the large (non-zero) entries of this singular vector reveals the position of the vertices of the clique. The algorithm computes the leading singular vector of $\tilde{A}$, and we are going to show that the large entries of this vector, with high probability, still correspond to the vertices of the clique.

From Theorem 1.2 it is easy to see that with probability at least .99, the ℓ_2 error $\|\tilde{\mathbf{u}}_1 - \mathbf{u}_1\|_2$ is bounded by $O(\sqrt{n}/k)$. Results from random matrix theory show that $\|E\|$ is at most $3\sqrt{n}$ [59], with probability $1 - o(1)$. Thus, our Theorem 3.3 implies that with probability at least .99, the infinity norm bound between the first singular vector of A and that of $\tilde{A}$ is

$$O(k^{-1/2}\sqrt{n}/k) \leq k^{-1/2}/4,$$

given that $k/\sqrt{n}$ is sufficiently large (beating the hidden constant in the big O). Thus, in the leading singular vector, the entries from the clique are at least $\frac{3}{4}k^{-1/2}$, and the rest are at most $\frac{1}{4}k^{-1/2}$ in absolute value. This guarantees that the clustering described in the algorithm reveals all the vertices of X . $\square$

While we have made no effort to optimize the value of c_0 (indeed, our theoretical constant is quite large), it is an interesting question to determine the values of c_0 for which FSC can recover the hidden clique exactly. The optimal c_0 is quite small; it is likely close to 1. See Figure 2.

REMARK 14.1. *The density $1/2$ is not critical, and can be replaced by any parameter $p > n^{-1+\epsilon}$ (or even $p > n^{-1}\log^c n$, for some properly chosen c). In the case of p , one needs to replace a zero entry by $-p/(1-p)$. The random matrix E now has zero mean and spectral norm at most $3\sqrt{np}$; see again [59]. Thus, by following our argument, we can see that it is sufficient to assume $k \geq C\sqrt{np}$ for a sufficiently large constant C .*

14.2 Clique partition. Let us consider the situation where one hides many cliques $X_1, \dots, X_r$ of size $k_1 \geq k_2 \geq \dots \geq k_r$, which form a partition of the vertex set. The vertices of different cliques are connected with probability $1/2$, independently. The first task is to find the i th largest clique X_i , for any given $1 \leq i \leq r$, given one instance of the random graph.

One can do this by finding all X_i and then sorting them out. However, we can do the task directly by just computing the i th singular vector and clustering on its coordinates (the same way as in the last section). Before starting the main step of the algorithm, we change all zeros in the adjacency matrix to -1 .




Figure 2: Numerical results for FSC on the Hidden Clique problem. For each c , 40 trials of the hidden clique problem are run with $n = 1000$. The step of switching 0 to -1 is omitted. The random graph in which the clique is embedded is $G(n, 1/2)$. Reported are the fraction of trials for which the clique was recovered exactly with no mis-classifications of any vertex.

ALGORITHM 14.2. (*i*TH CLIQUE) Compute the i th singular vector. Let x be the largest value of a coordinate and let X_i be the set of all coordinates with value at least $x/2$.

One issue is that if $k_i = k_{i-1}$, then there is no way to differentiate X_i from X_{i-1} . Thus, the hard instances for the problem are when $|k_i - k_j|$ are small in general. In what follows, we concentrate on that case, and assume that all k_i are of order n . Our theorems enable us to find X_i correctly under the assumption that $|k_i - k_{i\pm 1}| = \tilde{O}(1)$.

THEOREM 14.2. *For any constant $c > 0$ there is a constant C such that the following holds. Assume that $k_r \geq cn$ and $k_i - k_{i+1} \geq C \log n$ for all $1 \leq i \leq r$. Then with probability at least .99, Algorithm 14.2 recovers X_i correctly, for any $1 \leq i \leq r$.*

In what follows, we illustrate the ideas through the case $i = 1$. The analysis for a general i is similar. Consider the leading eigenvector of $\tilde{A} = A + E$, where A now consists of r disjoint diagonal all-one blocks of sizes $k_1, \dots, k_r$.

$$A = \begin{bmatrix} \mathbf{1}_{k_1} \mathbf{1}_{k_1}^T & 0 & \dots & 0 \\ 0 & \mathbf{1}_{k_2} \mathbf{1}_{k_2}^T & 0 & \vdots \\ \vdots & 0 & \ddots & 0 \\ 0 & \dots & 0 & \mathbf{1}_{k_r} \mathbf{1}_{k_r}^T \end{bmatrix}$$

The leading eigenvalue of A is k_1 and the leading eigenvector is

$$(1/\sqrt{k_1}, \dots, 1/\sqrt{k_1}, 0, \dots, 0).$$

The next eigenvalue is k_2 and the gap $\delta_1 = k_1 - k_2$. Since the cliques partition the vertex set, $k_1 \geq n/r$. The difference (compared to the previous section) is that, with probability .99, the ℓ_2 error (from Theorem 1.2) is now

bounded by

$$O\left[\frac{1}{\delta_1} + \frac{\|E\|}{\sigma} + \frac{\|E\|^2}{\delta_1 \sigma}\right] = O\left[\frac{1}{k_1 - k_2} + \frac{\sqrt{n}}{k_1} + \frac{n}{k_1(k_1 - k_2)}\right] = O\left[\frac{1}{k_1 - k_2}\right].$$

Since we assume that all $k_i = \Theta(n)$, the infinity norm of U is $O(\frac{1}{\sqrt{n}})$. So, our Theorem 3.3 implies that with probability at least .99, the infinity norm bound for the first eigenvector is

$$O\left[\frac{r\sqrt{\log n}}{(k_1 - k_2)\sqrt{n}} + \frac{\log n}{n}\right] \leq \frac{1}{4k_1^{1/2}},$$

given that $(k_1 - k_2)/\sqrt{\log n}$ is bounded from below by a sufficiently large constant, proving the claim.

We found a simple, but effective, trick to reduce the general case (when the separation condition could be violated, such as when k_i are all the same) to the situation in Theorem 14.2. We call this trick *random truncation* and it works as follows.

Random truncation. Select each vertex with probability $\rho := n^{-1+\epsilon}$, independently, where ϵ is a small positive constant. Let S be the set of selected vertices and $V' = V \setminus S$, $X'_i = X_i \setminus S$, $k'_i = |X'_i|$. If $k_i = \Theta(n)$ then $|X_i \cap S|$ is a binomial random variable χ with mean $k_i \rho = \Theta(n\rho)$ and standard deviation $\Theta(\sqrt{n\rho}) = \Theta(n^{\epsilon/2})$. Since $\log n = o(n^{\epsilon/2})$, the following fact is obvious.

FACT 14.1. (SEPARATION LEMMA) *Consider the random variable χ above and let χ' be its independent copy. Then for any given interval I of length $O(\log n)$, with probability at least $1 - o(1)$, $\chi - \chi' \notin I$.*

By the union bound over all pairs (i, j) , it follows that with probability at least $1 - o(1)$, $\min_{i \neq j} |k'_i - k'_j| = \omega(\log n)$. Thus, our separation condition holds on the subgraph spanned by V' . We can now run our algorithm on the adjacency matrix of this graph to identify X'_i . To finish, define X_i as the union of X'_i with the vertices in S which are connected to all the vertices in X'_i .

Another natural task is to find all X_i , and we can complete this task by consecutive applications of the algorithm FSC from the last section. First, find X_1 (or more precisely X'_1), then remove it from the graph. Then, find X_2 and continue in this way. Here is the formal description of the algorithm.

ALGORITHM 14.3. (CLIQUE PARTITION)

1. Define a set S by choosing each vertex in S with probability $\rho := n^{-1+\epsilon}$. Let $V' = V \setminus S$ and consider the graph spanned by V' .
2. For $i = 1, \dots, r-1$, run FSC to get X'_i . Let $X'_r = V' \setminus \cup_{i=1}^{r-1} V_i$.
3. Define X_i be the union of X'_i and the vertices of S which are adjacent to all of X'_i .

THEOREM 14.3. *Assume that $k_i = \Theta(n)$ for all $1 \leq i \leq r$. With probability at least .9, the Algorithm Clique Partition recovers $X_1, \dots, X_n$ correctly.*

14.3 Planted colorings. Finding a r coloring of a graph is a notoriously hard problem, even when we know that the graph is r -colorable. A number of researchers have considered the random instance of this problem. One natural setting is as follows. Partition the vertex set V into r independent sets $X_1, \dots, X_r$ of sizes $k_1 \geq k_2 \geq \dots \geq k_r$ and then connect the vertices between different X_i with probability $1/2$. The task is to recover the proper coloring from one instance of this random graph; see for instance [6], [13].

Notice that if we look at the complement graph, then this is exactly the problem considered in the previous section, as independent sets become cliques. Thus, we obtain

THEOREM 14.4. *Assume that $k_i = \Theta(n)$ for all $1 \leq i \leq r$. Then with probability at least .9, the algorithm in the last section recovers the planted coloring.*

REMARK 14.2. *In this and the previous problems, the constant $1/2$ again is not important, and can be replaced by a general density p . The condition that $k_i = \Theta(n)$ for all i can also be weakened.*

14.4 Hidden partition. We now consider a generalization of the problem in Section 14.2, where each clique X_i is replaced by a random graph with edge density $p_i > 1/2$. Similar to Section 14.2, the task is to locate a particular X_i or all X_i from one random instance of the graph.

Switch all 0 in the adjacency matrix to -1 . The resulting matrix $\tilde{A}$ can be decomposed into $A + E$, where A now has the following form

$$A = \begin{bmatrix} (2p_1 - 1)\mathbf{1}_{k_1}\mathbf{1}_{k_1}^T & 0 & \dots & 0 \\ 0 & (2p_2 - 1)\mathbf{1}_{k_2}\mathbf{1}_{k_2}^T & 0 & \vdots \\ \vdots & 0 & \ddots & 0 \\ 0 & \dots & 0 & (2p_r - 1)\mathbf{1}_{k_r}\mathbf{1}_{k_r}^T \end{bmatrix}$$

The random matrix E has the following form. The entry e_{ij} , $1 \leq i \leq j \leq n$, is Rademacher (± 1) if i and j belong to different X_k . If they belong to the same X_k , then let $e_{ij} = (2 - 2p_k)$ with probability p_k and $-2p_k$ with probability $1 - p_k$. It is easy to check that $\tilde{A} = A + E$ and that all entries of E have zero-mean and are 2-bounded.

Set $\rho_i := 2p_i - 1$. The singular values of A are $k_1\rho_1, \dots, k_r\rho_r$. We replace the assumption $k_1 \geq k_2 \geq \dots \geq k_r$ by its weighted version $k_1\rho_1 \geq k_2\rho_2 \geq \dots \geq k_r\rho_r$. The leading singular value of A is now $k_1\rho_1$, the second is $k_2\rho_2$ and the gap is $\delta_1 = k_1\rho_1 - k_2\rho_2$; the singular vector remains the same. If we follow the proof of Theorem 14.2, then the condition becomes

$$\frac{1}{k_i^{1/2}} \frac{r}{\delta_i} + \frac{\sqrt{\log n}}{k_i\rho_1} \leq C^{-1} \frac{1}{k_i^{1/2}}$$

for some sufficiently large constant C . This is equivalent to assuming that both δ_i and $\sqrt{k_i\rho_i}/\sqrt{\log n}$ are lower bounded by some sufficiently large constant C . The second one is equivalent to $\rho_i \geq C\sqrt{\log n/n}$ for some sufficiently large constant C .

THEOREM 14.5. *For any constant c_1 , there is a constant C such that the following holds. Assume that $k_1\rho_1/k_r\rho_r$ is bounded from above by c_1 . If $\delta_i \geq C \log n$ and $\rho_i \geq C\sqrt{\log n/n}$ then the i th clique algorithm recovers X_i correctly with probability .9.*

We can again apply the random truncation trick at the beginning to guarantee the separation condition. However, the application of this trick on this more general setting is slightly more technical than in the case of cliques, since it is less obvious how to assign the vertices from S to X'_i . To decide which X'_i a vertex $v \in S$ belongs to, we first choose a subset $Y_i \subset X'_i$ so that all Y_i has the same size cn for some constant $c > 0$. (This is doable because we assume that all X_i have size $\Theta(n)$.) Let i be the index where v has the most edges connected to Y_i and then add v to X'_i .

Notice that if $v \notin X_i$, then the number of edges between v and Y_i has distribution $\chi_0 = \text{Binom}(.5, cn)$. If $v \in X_i$ then it has distribution $\chi_i = \text{Binom}(p_i, cn)$. If $p_i - .5 := \rho_i/2 > C_0\sqrt{n \log n}$ for a sufficiently large constant C_0 (which may depend on c and r), then with probability at least .99, $\chi_i \geq \chi_0$ for all $v \in S$ and $1 \leq i \leq r$. This leads us to the following algorithm and theorem.

ALGORITHM 14.4. (HIDDEN PARTITION)

1. Define a set S by choosing each vertex in S with probability $\rho := n^{-1+\epsilon}$. Let $V' = V \setminus S$ and consider the graph spanned by V' .
2. For $i = 1, \dots, r-1$, run FSC to get X'_i . Let $X'_r = V' \setminus \cup_{i=1}^{r-1} V_i$.
3. Select subsets $Y_i \subset X'_i$ such that $|Y_i| = cn$, for some properly chosen small constant $c > 0$.
4. Define X_i be the union of X'_i and those vertices v of S where $d_i(v) = \max_j d_j(v)$, where $d_i(v)$ is the number of neighbors of v in Y_i (break ties arbitrarily).

THEOREM 14.6. *For any constant c there are constants c_0, C such that the following holds. If $k_i \geq cn$ and $\rho_i \geq C\sqrt{\log n/n}$ for all $1 \leq i \leq r$, then Algorithm Hidden Partition recovers all X_i correctly with probability .9.*

REMARK 14.3. The density $1/2$ again is not important, and can be replaced by a general density q .

Another well-known instance of this problem is the hidden bipartition problem. In this problem, $r = 2$ and the vertex set is partitioned into two sets of equal size $n/2$. Draw edges with probability p inside X_i and $q < p$ between X_1 and X_2 . The task is to recover the partition from one instance of the random graph. This particular case has been studied heavily by many researchers through 4 decades; see Table 2.

Bui, Chaudhuri, Leighton, Sipser '84 [16]	min-cut method	$p = \Omega(1/n), q = o(n^{-1-4/((p+q)n)})$
Dyer, Frieze '89 [35]	min-cut via degrees	$p - q = \Omega(1)$
Boppana '87 [15]	spectral method	$(p - q)/\sqrt{p + q} = \Omega(\sqrt{\log(n)/n})$
Snijders, Nowicki '97 [56]	EM algorithm	$p - q = \Omega(1)$
Jerrum, Sorkin '98 [42]	Metropolis algorithm	$p - q = \Omega(n^{-1/6+\epsilon})$
Condon, Karp '99 [28]	augmentation algorithm	$p - q = \Omega(n^{-1/2+\epsilon})$
Carson, Impagliazzo '01 [14]	hill-climbing algorithm	$p - q = \Omega(n^{-1/2} \log^4 n)$
Mcscherry '01 [50]	spectral method	$(p - q)/\sqrt{p} \geq \Omega(\sqrt{\log(n)/n})$
Bickel, Chen '09 [12]	N-G modularity	$(p - q)/\sqrt{p + q} = \Omega(\log(n)/\sqrt{n})$
Rohe, Chatterjee, Yu '11 [55]	spectral method	$p - q = \Omega(1)$
Abbe, Bandeira, Hull '14 [2]	maximum likelihood	$p - q = \Omega(\sqrt{\log(n)/n})$
Vu '18 [61]	spectral method	$(p - q)/p^{1/2} = \Omega(\sqrt{\log(n)/n})$
Abbe, Fan, Wang, Zhong '19 [4]	spectral method	$p - q = \Omega(\sqrt{\log(n)/n})$

Table 2: A recreation of the table in [2] surveying the hidden bipartition problem, with some recent additions.

In this case, our method (with some obvious modifications to replace $1/2$ by q) gives

THEOREM 14.7. If $p, q = \Theta(1)$, $p - q \geq C\sqrt{\log n/n}$ for all $1 \leq i \leq r$, then Algorithm Hidden Partition recovers all X_i correctly with probability .9.

The lower bound $\sqrt{\log n/n}$ is the current best on this problem; see [2]. From our analysis, it is clear that the same conclusion holds for equal partitions with any number of parts (more than 2). The condition $p, q = \Theta(1)$ can also be improved, and with p tending to zero, it becomes $(p - q)/\sqrt{p} \geq C\sqrt{\log n/n}$.

15 Application: Exact Matrix Completion from Few Entries

In this section, we prove Theorem 4.2. Let us recall that A is an integer matrix with rank $r = O(1)$, with entries bounded by an absolute constant (so $\|A\|_\infty = O(1)$). Let S be the sampled version of A where each entry is sampled (independently) with probability p , and the un-sampled entries are zeroed out. Then $\tilde{A} := p^{-1}S$ is an unbiased estimate of A . We analyze the simple spectral algorithm in Section 15 to recover A exactly from $\tilde{A}$. Set $W = [U, V]$, the concatenated matrix of U and V , where U and V are the left and right singular vectors of A respectively. Recall that we set $s := \max_i \{i : \sigma_i \geq \frac{1}{16r} \|W\|_\infty^{-2}\}$, $\bar{\delta} = \inf_{i \leq s} \delta_i$, and $\tilde{s} = \max_i \{i : \tilde{\sigma}_i \geq \frac{1}{8r} \|W\|_\infty^{-2}\}$. Let $a = \max\{\sqrt{\|A\|_\infty}, 2\}$. For concreteness, set c to be the constant $(a + 1) \times 2^{18} \times 7r^3$.

By the rounding step, in order to have exact recovery, we need to show that

$$\|A - B\|_\infty < \frac{1}{2}.$$

Let us consider the entry A_{12} . By the singular value decompositions of A and B ,

$$A_{12} = \sum_{i=1}^r \sigma_i u_{i1} v_{i2}, \quad B_{12} = \sum_{i=1}^{\tilde{s}} \tilde{\sigma}_i \tilde{u}_{i1} \tilde{v}_{i2}.$$

Let $A_{\tilde{s}}$ be the best rank $\tilde{s}$ approximation of A . Thus, $A = A_{\tilde{s}} + \sum_{i=\tilde{s}+1}^r \sigma_i u_i u_i^T$. Write $A_{\tilde{s}ij}$ for the ij entry of $A_{\tilde{s}}$. By the triangle inequality, we have

$$(15.145) \quad \begin{aligned} |B_{12} - A_{12}| &\leq |B_{12} - A_{\tilde{s}12}| + \left| \sum_{i=\tilde{s}+1}^r \sigma_i u_{i1} v_{i2} \right| \\ &= \left| \sum_{i=1}^{\tilde{s}} \tilde{\sigma}_i \tilde{u}_{i1} \tilde{v}_{i2} - \sum_{i=1}^{\tilde{s}} \sigma_i u_{i1} v_{i2} \right| + \left| \sum_{i=\tilde{s}+1}^r \sigma_i u_{i1} v_{i2} \right|. \end{aligned}$$

Bounding the second term on the RHS is fairly straightforward, and is handled in Lemma 15.2. The choice of the threshold $\tilde{s}$ plays an important role. Before stating the lemma, we will need the following tail bound on the norm of E . Its proof relies on a result from 10 and is given in Appendix D. This lemma will be used throughout this section.

LEMMA 15.1. *There exists an absolute constant C such that*

$$(15.146) \quad \mathbb{P}\{\|E\| \geq C\sqrt{Np^{-1}}\} \leq N^{-3}.$$

For convenience, we define

$$\mathcal{E} := \{\|E\| \leq C\sqrt{Np^{-1}}\},$$

where C is the constant above.

LEMMA 15.2. *With probability at least $1 - N^{-3}$,*

$$\left| \sum_{i=\tilde{s}+1}^r \sigma_i u_{i1} v_{i2} \right| \leq \frac{1}{4}.$$

In the next lemma, we bound the first term on the RHS of (15.145). This is where we use our refined bounds for the large K case. As we have mentioned, unlike the clustering problem, K can be quite large in the matrix completion setting.

LEMMA 15.3. *Under the conditions of Theorem 4.2, with probability at least $1 - 3N^{-2}$,*

$$|B_{12} - A_{\tilde{s}12}| < \frac{1}{4}.$$

Lemmas 15.2 and 15.3 applied to the RHS of (15.145) are enough to conclude the result of Theorem 4.2, because then the error will be strictly less than $\frac{1}{4} + \frac{1}{4} = \frac{1}{2}$, and will thus be rounded away. The choice of considering entry $(1, 2)$ is arbitrary; with probability $1 - N^{-1}$, the bound holds for all of the entries. The remainder of the section is dedicated to proving Lemmas 15.2 and 15.3.

Proof of Lemma 15.2. Suppose that $\mathcal{E}$, which has probability $1 - N^{-3}$ by Lemma 15.1, occurs. Recall that $\tilde{s} = \max_i \{i : \tilde{\sigma}_i \geq \frac{\|W\|_{\infty}^{-2}}{8r}\}$. Therefore, for all $i \geq \tilde{s} + 1$, it must be the case that $\tilde{\sigma}_i \leq \frac{1}{8r} \|W\|_{\infty}^{-2}$. By Fact 5.1 and Lemma 15.1, for $i \geq \tilde{s} + 1$,

$$\sigma_i \leq \tilde{\sigma}_i + \|E\| \leq \frac{\|W\|_{\infty}^2}{8r} + C\sqrt{Np^{-1}}.$$

By the density assumption, $p > N^{-1} \log^{4.03} N$. By the incoherence assumption, $\|W\|_{\infty} \leq cN^{-1/2}$. Therefore, $\sqrt{Np^{-1}} = o\left(\|W\|_{\infty}^{-2}\right)$. It follows that for $i \geq \tilde{s} + 1$,

$$\sigma_i \leq \frac{\|W\|_{\infty}^2}{4r}.$$

We have

$$\begin{aligned}
 (15.147) \quad \left| \sum_{i=\tilde{s}+1}^r \sigma_i u_{i1} v_{i2} \right| &\leq \sum_{i=\tilde{s}+1}^r \sigma_i |u_{i1}| |v_{i2}| \\
 &\leq \frac{\|W\|_\infty^{-2}}{4r} \sum_{i=\tilde{s}+1}^r |u_{i1}| |v_{i2}|.
 \end{aligned}$$

Because both $|u_{i1}| \leq \|W\|_\infty$ and $|v_{il}| \leq \|W\|_\infty$, it follows that

$$(15.148) \quad \left| \sum_{i=\tilde{s}+1}^r \sigma_i u_{i1} v_{i2} \right| \leq \frac{1}{4r} \sum_{i=\tilde{s}+1}^r 1 \leq \frac{1}{4}.$$

□

In order to establish Lemma 15.3, we need the following proposition and lemma.

PROPOSITION 15.1. *Suppose that $\mathcal{E}$ occurs. Then, $\tilde{s} \leq s$.*

Proof. Suppose $\mathcal{E}$ occurs. Assume towards contradiction that $\tilde{s} > s$. By definition of $\tilde{s}$, this means that there are at least $s+1$ singular values of $\tilde{A}$ larger than $\frac{\|W\|_\infty^{-2}}{8r}$, so

$$(15.149) \quad \tilde{\sigma}_{s+1} \geq \frac{\|W\|_\infty^{-2}}{8r}.$$

By definition of s , $\sigma_{s+1} \leq \frac{\|W\|_\infty^{-2}}{16r}$. Therefore, by Fact 5.1,

$$\tilde{\sigma}_{s+1} \leq \sigma_{s+1} + \|E\| \leq \frac{\|W\|_\infty^{-2}}{16r} + C\sqrt{Np^{-1}} = \frac{\|W\|_\infty^{-2}}{16r} + o(\|W\|_\infty^{-2}) < \frac{\|W\|_\infty^{-2}}{8r}.$$

where we use our previous observation that $\sqrt{Np^{-1}} = o(\|W\|_\infty^{-2})$. This is a contradiction with (15.149). □

LEMMA 15.4. *Set $m_\infty(i) = \max\{\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_\infty, \|\tilde{\mathbf{v}}_i - \mathbf{v}_i\|_\infty\}$ and $m_2(i) = \max\{\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2, \|\tilde{\mathbf{v}}_i - \mathbf{v}_i\|_2\}$. Under the conditions of Theorem 4.2, there exists a constant c_0 such that with probability at least $1 - 2N^{-2}$,*

$$(15.150) \quad \sup_{i \leq s} m_\infty(i) \leq \frac{c_0 \|W\|_\infty}{\log N}.$$

Proof of Lemma 15.3 given Lemma 15.4. Recall that

$$|B_{12} - A_{\tilde{s}12}| = \left| \sum_{i=1}^{\tilde{s}} \tilde{\sigma}_i \tilde{u}_{i1} \tilde{v}_{i2} - \sum_{i=1}^{\tilde{s}} \sigma_i u_{i1} v_{i2} \right|.$$

Letting $\tilde{u}_{i1} = u_{i1} + \Delta u_{i1}$ and $\tilde{v}_{i2} = v_{i2} + \Delta v_{i2}$, we have

$$(15.151) \quad B_{12} = \sum_{i=1}^{\tilde{s}} \tilde{\sigma}_i [u_{i1} v_{i2} + u_{i1} \Delta v_{i2} + v_{i2} \Delta u_{i1} + \Delta u_{i1} \Delta v_{i2}].$$

Let c_0 be the constant from Lemma 15.4. Suppose $\mathcal{E}$ and $\{\sup_{i \leq s} m_\infty(i) \leq \frac{c_0}{\log N}\}$ both occur. By the lemma and the union bound, this happens with probability at least $1 - 3N^{-2}$. Since $\mathcal{E}$ occurs, $\tilde{s} \leq s$ by Proposition 15.1. It follows that

$$(15.152) \quad |B_{12} - A_{\tilde{s}12}| \leq \sum_{i=1}^s |\tilde{\sigma}_i - \sigma_i| |u_{i1} v_{i2}| + d,$$

where

$$d := \sum_{i=1}^s \tilde{\sigma}_i [|u_{i1}| |\Delta v_{i2}| + |v_{i2}| |\Delta u_{i1}| + |\Delta u_{i1}| |\Delta v_{i2}|].$$

Observe that both $|\Delta u_{i1}|$ and $|\Delta v_{i2}|$ are bounded by $m_\infty(i)$. We consider the two terms on the RHS of (15.152) separately. For the first term, using Fact 5.1, $|\tilde{\sigma}_i - \sigma_i| \leq \|E\|$. Since $|u_{i1}| |v_{i2}| \leq \|W\|_\infty^2$, we have

$$(15.153) \quad \sum_{i=1}^s |\tilde{\sigma}_i - \sigma_i| |u_{i1} v_{i2}| \leq r \|E\| \|W\|_\infty^2$$

By the incoherence assumption, $\|W\|_\infty = O(N^{-1/2})$ so $\|W\|^2 = O(N^{-1})$. Since $\mathcal{E}$ occurs,

$$\|E\| \leq C \sqrt{N p^{-1}} = o(N),$$

where the last equality uses the assumed lower bound for p , $p > \frac{\log^{4.03} N}{N}$. It follows that $\|E\| \|W\|_\infty^2 = o(1)$.

$$(15.154) \quad \sum_{i=1}^s |\tilde{\sigma}_i - \sigma_i| |u_{i1} v_{i2}| = o(1).$$

Moving to the term d on the RHS of (15.152), we bound for $i \leq s$, $\tilde{\sigma}_i \leq \sigma_i + \|E\|$ by Fact 5.1. By the signal-to-noise condition and the fact that $\mathcal{E}$ occurs, $\sigma_s > \|E\|$. Thus, $\sigma_i + \|E\| \leq 2\sigma_i$. Therefore,

$$(15.155) \quad \begin{aligned} d &= \sum_{i=1}^s \tilde{\sigma}_i [|u_{i1}| |\Delta v_{i2}| + |v_{i2}| |\Delta u_{i1}| + |\Delta u_{i1}| |\Delta v_{i2}|] \\ &\leq 2 \sum_{i=1}^s \sigma_i [|u_{i1}| |\Delta v_{i2}| + |v_{i2}| |\Delta u_{i1}| + |\Delta u_{i1}| |\Delta v_{i2}|]. \end{aligned}$$

Then, using the bounds $|\Delta u_{i1}|, |\Delta v_{i1}| \leq m_\infty(i)$ and $|u_{i1}|, |v_{i1}| \leq \|W\|_\infty$, we have

$$(15.156) \quad d \leq 2 \sum_{i=1}^s \sigma_i [2 \|W\|_\infty m_\infty(i) + m_\infty^2(i)].$$

Because A has rank $r = O(1)$ and has $O(1)$ bounded entries, $\sigma_1 = O(N)$. Since $\|W\|_\infty = O(N^{-1/2})$ by the incoherence assumption and $\sup_{i \leq s} m_\infty(i) = o(\|W\|_\infty)$,

$$(15.157) \quad d = O(\sigma_1 \|W\|_\infty \sup_{i \leq s} m_\infty(i)) = o(N \|W\|_\infty^2) = o(1).$$

(15.154) and (15.157) give that

$$(15.158) \quad |B_{12} - A_{s12}| = o(1) < \frac{1}{4},$$

for a large enough N . $\square$

In order to prove Lemma 15.4, first establish that we can apply our refined (large K case) results for the ℓ_∞ perturbation of singular vectors because the strong stability condition holds.

LEMMA 15.5. Recall that a is the absolute constant $a = \max\{\sqrt{\|A\|_\infty}, 2\}$. Under the conditions of Theorem 4.2, for all $i \leq s$, the singular values and gaps (σ_i, δ_i) are all $(\frac{c}{a+1}, N^{-3}, 2)$ strongly stable under E .

Proof of Lemma 15.4 given Lemma 15.5. Since the entries of A are $O(1)$, $K = O(p^{-1})$. By Lemma 15.5 and the choice of c , the conditions for Theorem 3.6 to bound $m_\infty(i)$ hold with $\frac{c}{a+1}$, $\tau = N^{-3}$, $\nu = 2$, and $K = O(p^{-1})$. Applying this for all $i \leq s$, using the fact that $\bar{\delta}$ is the minimum δ_i among the first s singular values,

$$(15.159) \quad \sup_{i \leq s} m_\infty(i) = O \left(\kappa_s \|W\|_\infty \left[\sup_{i \leq s} m_2(i) + \frac{\|E\|}{\sigma_s} + \frac{\sqrt{\log N}}{p\bar{\delta}} \right] + \frac{\kappa_s \sqrt{p^{-1}N} \|W\|_\infty \log N}{2\sigma_s} \right),$$

with probability at least $1 - N^{-2}$. Suppose in addition $\mathcal{E}$ holds, which happens with probability at least $1 - N^{-3}$.

Recall that we wish to show that the LHS is less than $c \|W\|_\infty \log^{-1} N$. Let us start by bounding κ_s . We have shown that $\sigma_1 = O(N)$. By definition of s , $\sigma_s \geq \frac{\|W\|_\infty^{-2}}{16r} = \Omega(N)$, where we also use the incoherence assumption. It follows that $\kappa_s = O(1)$. To handle the last term in the RHS of (15.159), we use that $\frac{\sqrt{Np^{-1}}}{\sigma_s} = O(\log^{-2.01} N)$ by the signal-to-noise assumption, so

$$(15.160) \quad \frac{\kappa_s \sqrt{p^{-1}N} \|W\|_\infty \log N}{2\sigma_s} = o(\|W\|_\infty \log^{-1} N).$$

For the remaining term, we first bound $\sup_{i \leq s} m_2(i)$. We appeal to the ℓ_2 perturbation bounds of [51]. We adapt their results to our situation in Corollary B.1 in Appendix B which we can apply because $\bar{\delta} > cp^{-1} \log N$. This result gives that with probability at least $1 - N^{-3}$,

$$(15.161) \quad \sup_{i \leq s} m_2(i) = O \left[\frac{\sqrt{\log N}}{p\bar{\delta}} + \frac{\|E\|}{\sigma_s} + \frac{\|E\|^2}{\sigma_s \bar{\delta}} \right].$$

Using this bound and (15.160), (15.159) becomes

$$(15.162) \quad \sup_{i \leq s} m_\infty(i) = O \left(\kappa_s \|W\|_\infty \left[\frac{\sqrt{\log N}}{p\bar{\delta}} + \frac{\|E\|}{\sigma_s} + \frac{\|E\|^2}{\sigma_s \bar{\delta}} \right] \right) + o(\|W\|_\infty \log^{-1} N).$$

Since $\kappa_s = O(1)$, the proof of the lemma will be complete once we establish that there is a constant c_1 such that the sum of the three terms in the brackets is at most $O(c_1 \log^{-1} N)$. Let us start with the third term in the square brackets, $\frac{\|E\|^2}{\sigma_s \bar{\delta}}$. Since $\mathcal{E}$ holds,

$$\|E\|^2 = O(Np^{-1}).$$

By the gap assumption, $\bar{\delta}^{-1}p^{-1} = O(\frac{1}{\log N})$. Thus, since $\sigma_s = \Omega(N)$,

$$\|E\|^2 \bar{\delta}^{-1} \sigma_s^{-1} = O(N \sigma_s^{-1} \bar{\delta}^{-1} p^{-1}) = O(\bar{\delta}^{-1} p^{-1}) = O(\log^{-1} N).$$

Moving to the first term in the square brackets in (15.162), we have just established that $\bar{\delta}^{-1}p^{-1} = O(\log^{-1} N)$. Lastly, for the second term in the square brackets in (15.162), $\|E\| \sigma_s^{-1} = o(\log^{-1} N)$. This is by the signal-to-noise assumption and because $\mathcal{E}$ occurs. Since all three terms are either $O(\log^{-1} N)$ or $o(\log^{-1} N)$, the existence of c_1 can be quickly inferred.

A quick inspection of the proof gives that the total probability of occurrence of the events considered is at least $1 - 2N^{-2}$.

□

Proof of Lemma 15.5. Recall that $T = \inf\{t > 0 : \mathbb{P}(\|E\| > t) \leq N^{-1/3}\}$. By Lemma 15.1, $T \leq C\sqrt{Np^{-1}}$, where C is the absolute constant from the lemma. Recall that $\bar{\delta}$ is the smallest gap in the first s singular values of A . We will use σ_s and $\bar{\delta}$ to show that the singular values and gaps (σ_i, δ_i) for $i \leq s$ satisfy the $(\frac{c}{a+1}, N^{-3}, 2)$ strong stability condition. Let $1 \leq i \leq s$.

We first verify the three conditions for $(\frac{c}{a+1}, N^{-3}, 2)$ stability, and conclude with verifying strong stability. First, the signal-to-noise condition gives

$$(15.163) \quad \sigma_i \geq \sigma_s > c\sqrt{Np^{-1}} \log^{2.01} N,$$

which ensures that

$$\sigma_i \geq \sigma_s > c\sqrt{Np^{-1}} \log^{2.01} N > \frac{c}{2}T \geq \frac{c}{a+1}T.$$

This shows that condition (a) in Definition 3.1 holds for (σ_i, δ_i) with $\frac{c}{a+1}$ and $\tau = N^{-3}$.

Next, recall the gap condition

$$(15.164) \quad \bar{\delta} > cp^{-1} \log N.$$

In order to verify (c) of Definition 3.1, we will show that

$$(15.165) \quad \bar{\delta} > \frac{c}{2}T\kappa_i \|W\|_\infty,$$

which will imply $\delta_i > \frac{c}{a+1}T\kappa_i \|W\|_\infty$, as desired. Because $\kappa_s = O(1)$, the bound on T implies

$$T\kappa_i \|W\|_\infty = O(\sqrt{p^{-1}}\sqrt{N} \|W\|_\infty) = O(\sqrt{p^{-1}})$$

by the incoherence assumption. Thus, (15.165) holds by (15.164).

We examine (b) in Definition 3.1, the final condition to verify $(\frac{c}{a+1}, N^{-3}, 2)$ stability. Recall that K is the bound for the absolute value of the entries of E . Observe that $K \leq \|A\|_\infty p^{-1}$. Since $\nu = 2$, we must show that

$$(15.166) \quad \delta_i > \frac{c}{a+1} \left(K \log N + \frac{T^2}{\sigma_s} \right).$$

Since $T \leq C\sqrt{Np^{-1}}$, it is sufficient to show that

$$(15.167) \quad \bar{\delta} > \frac{c}{a+1} \left(K \log N + C^2 \frac{Np^{-1}}{\sigma_s} \right).$$

Equation (15.164) implies that

$$\bar{\delta} > cp^{-1} \log N = \frac{c}{\|A\|_\infty} \|A\|_\infty p^{-1} \log N \geq \frac{c}{\|A\|_\infty} K \log N \geq \frac{c}{a} K \log N.$$

Since $\sigma_s = \Omega(N)$, $\frac{C^2 Np^{-1}}{\sigma_s} = o(K \log N)$ because $p > N^{-1} \log N$. Therefore,

$$\bar{\delta} > \frac{c}{a} K \log N > \frac{c}{a+1} \left(K \log N + C^2 \frac{Np^{-1}}{\sigma_s} \right)$$

as desired.

Having established $(\frac{c}{a+1}, N^{-3}, 2)$ stability, we conclude with the verification of strong stability. The signal to noise condition, equation (15.163), implies that

$$(15.168) \quad \sigma_i \geq \sigma_s > \frac{c}{\sqrt{\|A\|_\infty}} \sqrt{N \|A\|_\infty p^{-1}} \log^{2.01} N \geq \frac{c}{\sqrt{\|A\|_\infty}} \sqrt{NK} \log^{2.01} N \geq \frac{c}{a+1} \sqrt{NK} \log^{2.01} N.$$

Thus, the conditions for strong stability in Definition 3.2 are satisfied for (σ_i, δ_i) with $\frac{c}{a+1}$, $\tau = N^{-3}$, and $\nu = 2$. $\square$

References

- [1] E. ABBE, *Community detection and stochastic block models: Recent developments*, Journal of Machine Learning Research, 18 (2018), pp. 1–86.
- [2] E. ABBE, A. S. BANDEIRA, AND G. HALL, *Exact recovery in the stochastic block model*, IEEE Transactions on Information Theory, 62 (2016), pp. 471–487.
- [3] E. ABBE, J. FAN, AND K. WANG, *An ℓ_p theory of pca and spectral clustering*, The Annals of Statistics, 50 (2022).
- [4] E. ABBE, J. FAN, K. WANG, AND Y. ZHONG, *Entrywise eigenvector analysis of random matrices with low expected rank*, Annals of Statistics, 48 (2020), pp. 1452–1474.
- [5] J. AGTERBERG, Z. LUBBERTS, AND C. E. PRIEBE, *Entrywise estimation of singular vectors of low-rank matrices with heteroskedasticity and dependence*, IEEE Transactions on Information Theory, 68 (2021), pp. 4618–4650.
- [6] N. ALON AND N. KAHALE, *A spectral technique for coloring random 3-colorable graphs*, in Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, STOC '94, New York, NY, USA, 1994, Association for Computing Machinery, p. 346–355.
- [7] N. ALON, M. KRIVELEVICH, AND B. SUDAKOV, *Finding a large hidden clique in a random graph*, in Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '98, USA, 1998, Society for Industrial and Applied Mathematics, p. 594–598.
- [8] N. ALON AND J. H. SPENCER, *The Probabilistic Method*, Wiley Publishing, 4th ed., 2016.
- [9] M. AVELLA-MEDINA, H. S. BATTEY, J. FAN, AND Q. LI, *Robust estimation of high-dimensional covariance and precision matrices*, Biometrika, 105 (2018), pp. 271–284.
- [10] A. S. BANDEIRA AND R. VAN HANDEL, *Sharp nonasymptotic bounds on the norm of random matrices with independent entries*, The Annals of Probability, 44 (2016), pp. 2479 – 2506.
- [11] Z. BAO, X. DING, AND K. WANG, *Singular vector and singular subspace distribution for the matrix denoising model*, The Annals of Statistics, 49 (2021).
- [12] P. J. BICKEL AND A. CHEN, *A nonparametric view of network models and newman–girvan and other modularities*, Proceedings of the National Academy of Sciences, 106 (2009), pp. 21068 – 21073.
- [13] A. BLUM AND J. SPENCER, *Coloring random and semi-random k -colorable graphs*, J. Algorithms, 19 (1995), p. 204–234.
- [14] B. BOLLOBÁS AND A. D. SCOTT, *MAX-CUT for random graphs with a planted partition*, Comb. Probab. Comput., 13 (2004), p. 451–474.
- [15] R. B. BOPPANA, *Eigenvalues and graph bisection: An average-case analysis*, in 28th Annual Symposium on Foundations of Computer Science (sfcs 1987), 1987, pp. 280–285.
- [16] T. BUI, S. CHAUDHURI, F. LEIGHTON, AND M. SIPSER, *Graph bisection algorithms with good average case behavior*, Combinatorica, 7 (1987), pp. 171–191.
- [17] C. CAI, G. LI, Y. CHI, H. V. POOR, AND Y. CHEN, *Subspace estimation from unbalanced and incomplete data matrices: $\ell_{2,\infty}$ statistical guarantees*, The Annals of Statistics, 49 (2021), pp. 944 – 967.
- [18] E. J. CANDÈS AND Y. PLAN, *Matrix completion with noise*, Proceedings of the IEEE, 98 (2009), pp. 925–936.
- [19] E. J. CANDÈS AND B. RECHT, *Exact matrix completion via convex optimization*, Commun. ACM, 55 (2012), p. 111–119.

- [20] E. J. CANDÈS AND T. TAO, *The power of convex relaxation: Near-optimal matrix completion*, IEEE Trans. Inf. Theor., 56 (2010), p. 2053–2080.
- [21] E. J. CANDÈS AND M. B. WAKIN, *An introduction to compressive sampling*, IEEE Signal Processing Magazine, 25 (2008), pp. 21–30.
- [22] J. CAPE, M. TANG, AND C. E. PRIEBE, *The two-to-infinity norm and singular subspace geometry with applications to high-dimensional statistics*, The Annals of Statistics, (2017).
- [23] J. CAPE, M. TANG, AND C. E. PRIEBE, *Signal-plus-noise matrix models: Eigenvector deviations and fluctuations*, Biometrika, 106 (2019), p. 243–250.
- [24] S. CHATTERJEE, *Matrix estimation by universal singular value thresholding*, The Annals of Statistics, 43 (2015).
- [25] Y. CHEN, Y. CHI, J. FAN, AND C. MA, *Spectral methods for data science: A statistical perspective*, Foundations and Trends in Machine Learning, 14 (2021), pp. 566–806.
- [26] Y. CHEN, Y. CHI, J. FAN, C. MA, AND Y. YAN, *Noisy matrix completion: Understanding statistical guarantees for convex relaxation via nonconvex optimization*, SIAM Journal on Optimization, 30 (2020), pp. 3098–3121.
- [27] C. CHENG, Y. WEI, AND Y. CHEN, *Tackling small eigen-gaps: Fine-grained eigenvector estimation and inference under heteroscedastic noise*, IEEE Trans. Inf. Theor., 67 (2021), p. 7380–7419.
- [28] A. CONDON AND R. M. KARP, *Algorithms for graph partitioning on the planted partition model*, in Random Struct. Algorithms, 1999.
- [29] C. DAVIS AND W. M. KAHAN, *The rotation of eigenvectors by a perturbation. iii*, SIAM Journal on Numerical Analysis, 7 (1970), pp. 1–46.
- [30] Y. DEKEL, O. GUREL-GUREVICH, AND Y. PERES, *Finding Hidden Cliques in Linear Time with High Probability*, pp. 67–75.
- [31] P. B. DENTON, S. J. PARKE, T. TAO, AND X. ZHANG, *Eigenvectors from eigenvalues: A survey of a basic identity in linear algebra*, Bulletin, new series, of the American Mathematical Society, 59 (2021).
- [32] P. B. DENTON, S. J. PARKE, AND X. ZHANG, *Neutrino oscillations in matter via eigenvalues*, Phys. Rev. D, 101 (2020), p. 093001.
- [33] L. DING AND Y. CHEN, *The leave-one-out approach for matrix completion: Primal and dual analysis*, IEEE Transactions on Information Theory, PP (2018).
- [34] P. DRINEAS, A. FRIEZE, R. KANNAN, S. VEMPALA, AND V. VINAY, *Clustering large graphs via the singular value decomposition: Theoretical advances in data clustering (guest editors: Nina mishra and rajeev motwani)*, Machine Learning, 56 (2004).
- [35] M. E. DYER AND A. M. FRIEZE, *The solution of some random NP-hard problems in polynomial expected time*, J. Algorithms, 10 (1989), p. 451–489.
- [36] J. ELDRIDGE, M. BELKIN, AND Y. WANG, *Unperturbed: spectral analysis beyond davis-kahan*, in Proceedings of Algorithmic Learning Theory, F. Janoos, M. Mohri, and K. Sridharan, eds., vol. 83 of Proceedings of Machine Learning Research, PMLR, 4 2018, pp. 321–358.
- [37] J. FAN, W. WANG, AND Y. ZHONG, *An ℓ_∞ eigenvector perturbation bound and its application to robust covariance estimation*, Journal of Machine Learning Research, 18 (2016).
- [38] U. FEIGE AND R. KRAUTHGAMER, *Finding and certifying a large hidden clique in a semirandom graph*, Random Struct. Algorithms, 16 (2000), p. 195–208.

- [39] U. FEIGE AND E. O. OFEK, *Spectral techniques applied to sparse random graphs*, Random Structures & Algorithms, 27 (2005).
- [40] U. FEIGE AND D. RON, *Finding hidden cliques in linear time*, Discrete Mathematics & Theoretical Computer Science, DMTCS Proceedings vol. AM, 21st International Meeting on Probabilistic, Combinatorial, and Asymptotic Methods in the Analysis of Algorithms (AofA'10) (2010).
- [41] W. FERNANDEZ DE LA VEGA, *MAX-CUT has a randomized approximation scheme in dense graphs*, Random Structures & Algorithms, 8 (1996), pp. 187–198.
- [42] M. JERRUM AND G. B. SORKIN, *The metropolis algorithm for graph bisection*, Discret. Appl. Math., 82 (1998), pp. 155–175.
- [43] R. KANNAN AND S. VEMPALA, *Spectral algorithms*, Foundations and Trends in Theoretical Computer Science, 4 (2009), pp. 157–288.
- [44] R. H. KESHAVAN, A. MONTANARI, AND S. OH, *Matrix completion from a few entries*, IEEE Transactions on Information Theory, 56 (2010), pp. 2980–2998.
- [45] R. H. KESHAVAN, A. MONTANARI, AND S. OH, *Matrix completion from noisy entries*, Journal of Machine Learning Research, 11 (2010), pp. 2057–2078.
- [46] V. KOLTCHINSKII AND D. XIA, *Perturbation of linear forms of singular vectors under gaussian noise*, High Dimensional Probability VII, (2016), p. 397–423.
- [47] Y. KOREN, R. BELL, AND C. VOLINSKY, *Matrix factorization techniques for recommender systems*, Computer, 42 (2009), pp. 30–37.
- [48] X. P. LI, L. HUANG, H. C. SO, AND B. ZHAO, *A survey on matrix completion: Perspective of signal processing*, arXiv: Signal Processing, (2019).
- [49] S. LING, *Near-optimal performance bounds for orthogonal and permutation group synchronization via spectral methods*, Applied and Computational Harmonic Analysis, 60 (2022), p. 20–52.
- [50] F. MCSHERRY, *Spectral partitioning of random graphs*, in Proceedings 42nd IEEE Symposium on Foundations of Computer Science, 2001, pp. 529–537.
- [51] S. O’ROURKE, V. VU, AND K. WANG, *Random perturbation of low rank matrices: Improving classical bounds*, Linear Algebra and its Applications, 540 (2017).
- [52] S. O’ROURKE, V. VU, AND K. WANG, *Matrices with gaussian noise: optimal estimates for singular subspace perturbation*, 2023.
- [53] S. PÉCHÉ, *The largest eigenvalue of small rank perturbations of hermitian random matrices*, Probability Theory and Related Fields, 134 (2006), pp. 127–173.
- [54] B. RECHT, *A simpler approach to matrix completion*, J. Mach. Learn. Res., 12 (2011), p. 3413–3430.
- [55] K. ROHE, S. CHATTERJEE, AND B. YU, *Spectral clustering and the high-dimensional stochastic blockmodel*, The Annals of Statistics, 39 (2011).
- [56] T. SNIJDERS AND K. NOWICKI, *Estimation and prediction for stochastic blockmodels for graphs with latent block structure*, Journal of Classification, 14 (1997), pp. 75–100.
- [57] M. UDELL AND A. TOWNSEND, *Why are big data matrices approximately low rank?*, SIAM Journal on Mathematics of Data Science, 1 (2019), pp. 144–160.
- [58] R. VERSHYNIN, *High-Dimensional Probability: An Introduction with Applications in Data Science*, no. 47 in Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press.

- [59] V. VU, *Spectral norm of random matrices*, *Combinatorica*, 27 (2005), pp. 721–736.
- [60] V. VU, *Singular vectors under random perturbation*, *Random Structures & Algorithms*, 39 (2011), pp. 526–538.
- [61] V. VU, *A simple SVD algorithm for finding hidden partitions*, *Combinatorics, Probability and Computing*, 27 (2018), p. 124–140.
- [62] P.-A. WEDIN, *Perturbation bounds in connection with singular value decomposition*, *BIT*, 12 (1972), p. 99–111.
- [63] A. Y. ZHANG AND H. H. ZHOU, *Leave-one-out singular subspace perturbation analysis for spectral clustering*, 2022.
- [64] Y. ZHONG, *Eigenvector under random perturbation: A nonasymptotic rayleigh-schrödinger theory*, 2017.

A Perturbation of Singular values

We begin with the definition of the *concentration property*. The authors of [51] state that a square matrix E satisfies the (C, c, γ) concentration property if for all unit vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$, and $t > 0$,

$$(A.1) \quad \mathbb{P}(|\mathbf{u}^T E \mathbf{v}| > t) \leq C \exp(-ct^\gamma).$$

By using Hoeffding's inequality, they show that if $K \geq 1$ and E is an $n \times n$ symmetric matrix with independent, K bounded entries, then

$$(A.2) \quad \mathbb{P}(|\mathbf{u}^T E \mathbf{v}| > t) \leq 2 \exp\left(-\frac{1}{8K^2}t^2\right)$$

In other words, E satisfies the $(2, \frac{1}{8K^2}, 2)$ -concentration property.

A key ingredient in the analysis is the singular values of the perturbed matrices $A^{\{l\}}$ and $\tilde{A}$. It is very easy to verify that if E is a symmetric random matrix with independent, K -bounded entries, then for all l , $E^{\{l\}}$ has the $(2, \frac{1}{8K^2}, 2)$ concentration property as well. Therefore, the results derived in [51] for the perturbation of the singular values applies to $A^{\{l\}}$ for all l .

The main result for the perturbation of singular values in [51] is the following theorem.

THEOREM A.1. *Suppose that E is (C, c, γ) concentrated. Suppose that A has rank r , and let $1 \leq i \leq r$ be an integer. Then, for any $t \geq 0$,*

$$(A.3) \quad \tilde{\sigma}_j \geq \sigma_j - t$$

with probability at least

$$(A.4) \quad 1 - 2C9^i \exp\left(-c\frac{t^\gamma}{4^\gamma}\right)$$

and

$$(A.5) \quad \tilde{\sigma}_i \leq \sigma_i + tr^{1/\gamma} + 2\sqrt{i} \frac{\|E\|^2}{\tilde{\sigma}_i} + i \frac{\|E\|^3}{\tilde{\sigma}_i^2}$$

with probability at least

$$(A.6) \quad 1 - 2C9^{2r} \exp\left(-cr\frac{t^\gamma}{4^\gamma}\right).$$

Theorem 5.1 follows from Theorem A.1 and the fact that a symmetric, K bounded random matrix with independent (above the diagonal) entries satisfies the $(2, \frac{1}{8K^2}, 2)$ concentration property.

B ℓ_2 Perturbation of Eigenvectors

THEOREM B.1. ([51]) *Suppose that E is (C, c, γ) concentrated for a trio of constants (C, c, γ) and suppose that A has rank r . Then, for any $t > 0$,*

$$(B.7) \quad \|\tilde{\mathbf{u}}_1 - \mathbf{u}_1\|_2 \leq 8\left(\frac{tr^{1/\gamma}}{\delta} + \frac{\|E\|}{\sigma_1} + \frac{\|E\|^2}{\delta\sigma_1}\right)$$

with probability at least

$$(B.8) \quad 1 - 54C \exp\left(-c\frac{\delta^\gamma}{8^\gamma}\right) - 2C9^{2r} \exp\left(-cr\frac{t^\gamma}{4^\gamma}\right).$$

Recall that a symmetric E with independent, mean zero, K bounded entries satisfies the $(2, \frac{1}{8K^2}, 2)$ -concentration property. We will be considering only such matrices E in what follows. Applying Theorem [B.1](#) gives the following theorem.

THEOREM B.2. *Let E be a random, K -bounded, symmetric matrix with independent entries above the diagonal. For any $t > 0$,*

$$(B.9) \quad \|\tilde{\mathbf{u}}_1 - \mathbf{u}_1\|_2 \leq 8 \left(\frac{tr^{1/2}}{\delta} + \frac{\|E\|}{\sigma_1} + \frac{\|E\|^2}{\delta\sigma_1} \right)$$

with probability at least

$$(B.10) \quad 1 - 108 \exp \left(-\frac{\delta^2}{8^3 K^2} \right) - 4 \times 9^{2r} \exp \left(-r \frac{t^2}{128 K^2} \right).$$

As a consequence, if $\frac{\delta}{K\sqrt{8^3}} = \tau$, then for all $t \geq 0$, with probability at least $1 - 108 \exp(-\tau^2) - 4 \times 9^{2r} \exp(-t)$,

$$(B.11) \quad \|\tilde{\mathbf{u}}_1 - \mathbf{u}_1\|_2 \leq 8 \left(\frac{K\sqrt{128}t}{\delta} + \frac{\|E\|}{\sigma_1} + \frac{\|E\|^2}{\delta\sigma_1} \right)$$

They also obtain the following recursive result for the perturbation of the remaining eigenvectors.

THEOREM B.3. *Assume that E is (C, c, γ) concentrated. Suppose that A has rank r , and let $1 \leq i \leq r$ be an integer. Then for any $t > 0$,*

$$(B.12) \quad \|\mathbf{u}_i - \tilde{\mathbf{u}}_i\|_2 \leq 16 \left[\sum_{j=1}^{i-1} \|\tilde{\mathbf{u}}_j - \mathbf{u}_j\|_2 + \frac{tr^{1/\gamma}}{\delta_j} + \frac{\|E\|}{\sigma_j} + \frac{\|E\|^2}{\sigma_j \delta_j} \right]$$

with probability at least

$$(B.13) \quad 1 - 6C9^i \exp \left(-c \frac{\delta_j^\gamma}{8^\gamma} \right) - 2C9^{2r} \exp \left(-cr \frac{t^\gamma}{4^\gamma} \right)$$

Fix $1 \leq s \leq r$. s represents a cutoff point beyond which the eigenvectors are not of interest. Let $\epsilon_j(t) := 16 \left(\frac{tr^{1/\gamma}}{\delta_j} + \frac{\|E\|}{\sigma_j} + \frac{\|E\|^2}{\sigma_j \delta_j} \right)$. By taking the union bound over the first s singular vectors and iterating this recursive bound, and letting $\bar{\delta}$ be the smallest gap in the first s singular values, we obtain

THEOREM B.4. *For all $i \leq s$, and $t > 0$,*

$$(B.14) \quad \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 \leq \sum_{j=1}^i \left[\sum_{k=0}^{i-j} 16^k \right] \epsilon_j(t)$$

with probability at least

$$(B.15) \quad 1 - 6sC9^s \exp \left(-c \frac{\bar{\delta}^\gamma}{8^\gamma} \right) - 2Cs9^{2r} \exp \left(-cr \frac{t^\gamma}{4^\gamma} \right).$$

As a consequence, letting $\tau = \frac{\bar{\delta}}{K\sqrt{8^3}}$ if E is symmetric and K -bounded with independent entries, there exists C depending only on s such that with probability at least $1 - C[\exp(-\tau^2) - 9^{2r} \exp(-t)]$,

$$(B.16) \quad \sup_{i \leq s} \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2 \leq C \left(\frac{Kt}{\bar{\delta}} + \frac{\|E\|}{\sigma_s} + \frac{\|E\|^2}{\bar{\delta}\sigma_s} \right),$$

for any $t > 0$. In the rectangular case, where A is an $m \times n$ matrix, and when $r = O(1)$, we derive the following corollary using the standard symmetrization trick. We encounter this setting in the matrix completion problem. Let $N = m + n$.

COROLLARY B.1. *If $r = O(1)$ and $\bar{\delta} > 400K\sqrt{\log N}$, with probability at least $1 - N^{-3}$, we have*

$$(B.17) \quad \sup_{i \leq s} \max\{\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_2, \|\tilde{\mathbf{v}}_i - \mathbf{v}_i\|_2\} = O\left(\frac{Kt}{\bar{\delta}} + \frac{\|E\|}{\sigma_s} + \frac{\|E\|^2}{\bar{\delta}\sigma_s}\right).$$

C Results for Section 13

LEMMA C.1. *Let $Y \in \overline{\mathcal{L}_{\{\cdot, l\}}}$ holds. Then,*

$$(C.18) \quad \mathbb{P}(\mathcal{I}'_l | E^{\{l\}} = Y) \leq 2 \exp(-c_2 \log n).$$

Proof of Proposition 13.1 given the lemma. The proposition follows immediately from the lemma once we observe that

$$\mathbb{P}(\mathcal{I}'_l \cap \overline{\mathcal{L}_{\{\cdot, l\}}}) \leq \sup_{Y \in \overline{\mathcal{L}_{\{\cdot, l\}}}} \mathbb{P}(\mathcal{I}'_l | E^{\{l\}} = Y) \leq 2 \exp(-c_2 \log n).$$

□

Proof of Lemma C.1. We are looking for a bound of

$$\mathbb{P}(\mathcal{I}'_l | E^{\{l\}} = Y) = \mathbb{P}\{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x}(l) \rangle| \geq c_2 \sqrt{2Kn} f_{j+1} \log n \mid E^{\{l\}} = Y\}.$$

The lemma only considers realizations Y of $E^{\{l\}}$ satisfying $\|\mathbf{u}_i^{\{l\}}\|_\infty \leq f_1$. Recall that $\mathbf{x}(l)$ is essentially the l th row of E and $l \in \beta$. Therefore, conditional on such $E^{\{l\}}$, $\mathbf{u}_i^{\{l\}}$ is a deterministic unit vector whose entries have absolute value at most f_1 . The only randomness in each event thus comes from $\mathbf{x}(l)$. It follows that the inner product $\langle \mathbf{u}_i^{\{l\}}, \mathbf{x}(l) \rangle$, conditional on $E^{\{l\}}$, is the sum of independent, Kf_1 bounded, mean zero random variables $x_k u_{ik}^{\{l\}}$. Thus, this quantity can be bounded with Bernstein's inequality. We are applying Bernstein's inequality conditionally, so we also need to find a bound for the sum of the conditional second moments of the $x_k u_{ik}^{\{l\}}$. Since $\mathbf{u}_i^{\{l\}}$ is deterministic when we condition on $E^{\{l\}}$, and $\mathbf{x}(l)$ is independent of $E^{\{l\}}$, we have

$$(C.19) \quad \sum_{k=1}^n \mathbb{E}[x_k^2 u_{ik}^{\{l\}2} | E^{\{l\}} = Y] = \sum_{k=1}^n u_{ik}^{\{l\}2} \mathbb{E}[x_k^2] \leq K \sum_{k=1}^n u_{ik}^{\{l\}2} = K.$$

For the inequality, we use that the second moments of the entries of E are at most K by Assumption 3.1. Applying Bernstein's inequality (Lemma 5.2) then gives us that

$$(C.20) \quad \begin{aligned} \mathbb{P}\left\{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x}(l) \rangle| > t \mid E^{\{l\}} = Y\right\} &\leq 2 \exp\left(\frac{-t^2/2}{\sum_{k=1}^n \mathbb{E}[x_k^2 u_{ik}^{\{l\}2} | E^{\{l\}} = Y] + Kf_1 t/3}\right) \\ &\leq 2 \exp\left(\frac{-t^2/2}{K + Kf_1 t/3}\right). \end{aligned}$$

Set $t = c_2 \sqrt{2Kn} f_1 \log n$. Since $n^{-1/2} \leq \|U\|_\infty \leq f_1$, we obtain that $K \leq Knf_1^2$. Thus, using this bound for the first term in the denominator of the RHS of (C.20),

$$\begin{aligned} \mathbb{P}\left\{|\langle \mathbf{u}_i^{\{l\}}, \mathbf{x}(l) \rangle| > c_2 \sqrt{2Kn} f_1 \log n \mid E^{\{l\}} = Y\right\} &\leq 2 \exp\left(\frac{-c_2^2 Kn f_1^2 \log^2 n}{Kn f_1^2 + \frac{\sqrt{2}c_2}{3} K f_1^2 \sqrt{Kn} \log n}\right) \\ (C.21) \qquad \qquad \qquad &\leq 2 \exp\left(\frac{-c_2^2 Kn f_1^2 \log^2 n}{Kn f_1^2 + \frac{\sqrt{2}c_2}{3} K f_1^2 n \log n}\right) \\ &\leq 2 \exp(-c_2 \log n). \end{aligned}$$

In the second line, we used that $K \leq n$ by Assumption 3.1 $\square$

D Proof of Lemma 15.1

In this section, we bound the spectral norm of the matrix E from the matrix completion problem. We will use the following result from [10], which is a tail bound for the norm of K bounded random matrices with independent entries.

THEOREM D.1. (REMARK 3.13 IN [10]) *Let X be a symmetric, mean zero $n \times n$ random matrix whose entries above the diagonal are independent. Suppose the entries of X , ξ_{ij} , are K bounded random variables. Let*

$$v = \max_i \sum_j \mathbb{E}[\xi_{ij}^2].$$

Then there exists a universal constant $c > 0$ such that for any $t \geq 0$,

$$\mathbb{P}\{\|E\| \geq 4\sqrt{v} + t\} \leq n \exp\left(-\frac{t^2}{cK^2}\right).$$

Proof of Lemma 15.1. Recall that E is an $m \times n$ random matrix with independent, K bounded entries, where $K = O(p^{-1})$. Form the symmetrization of E as

$$S = S(E) = \begin{bmatrix} 0 & E \\ E^T & 0 \end{bmatrix}.$$

S is a symmetric $N \times N$ random matrix satisfying the conditions of Theorem D.1. It is easy to check that $\|E\| \leq \|S\|$. Let the entries of S be given by s_{ij} . We need to calculate $v = \max_i \sum_j \mathbb{E}[s_{ij}^2]$ to apply the theorem. Recall that in the matrix completion setting, the entries of E have second moment at most K . It follows that $\mathbb{E}[s_{ij}^2] \leq K$, so $v \leq NK$.

By Theorem D.1

$$(D.22) \qquad \qquad \mathbb{P}\{\|S\| \geq 4\sqrt{NK} + t\} \leq N \exp\left(\frac{-t^2}{cK^2}\right).$$

Set $t = K\sqrt{4c \log N}$. It follows that

$$(D.23) \qquad \mathbb{P}\{\|S\| \geq 4\sqrt{NK} + K\sqrt{4c \log N}\} \leq N \exp(-4 \log N).$$

Since $\|E\| \leq \|S\|$,

$$(D.24) \qquad \mathbb{P}\{\|E\| \geq 4\sqrt{NK} + \sqrt{4cK \log N}\} \leq N^{-3}.$$

Recall that $K = O(p^{-1})$. Under the assumptions of Theorem 4.2, $p > \frac{\log N}{N}$, so $\sqrt{K \log N} = o(\sqrt{N})$. It follows that there exists an absolute constant C such that

$$(D.25) \qquad \mathbb{P}\{\|E\| \geq C\sqrt{NK}\} \leq N^{-3}.$$

$\square$